

Title	A comparative study of binary synthesis modelling software for neutron star mass estimates
Authors	Kennelly, Sorchá Riain
Publication date	2025
Original Citation	Kennelly, S. R. 2025. A comparative study of binary synthesis modelling software for neutron star mass estimates. MSc Thesis, University College Cork.
Type of publication	Masters thesis (Research)
Rights	© 2025, Sorchá Riain Kennelly. - https://creativecommons.org/publicdomain/zero/1.0/
Download date	2026-10-29 19:52:41
Item downloaded from	https://hdl.handle.net/10468/17965

A COMPARATIVE STUDY OF BINARY SYNTHESIS MODELLING SOFTWARE FOR NEUTRON STAR MASS ESTIMATES

Sorcha Riain Kennelly, B.Sc



UNIVERSITY COLLEGE CORK
SCHOOL OF PHYSICS

Thesis submitted for the degree of
Master of Science

March 2025

Supervisor: Dr. Mark Kennedy
Head of School: Prof. Paul Callanan

Contents


1	Introduction	1
1.1	Neutron Stars	1
1.2	Spider Pulsars	2
1.3	Spiders and the EOS	5
1.4	The Mass Gap Question	7
2	Binary Modelling	9
2.1	Photometry	10
2.2	Modelling Codes	11
2.2.1	Icarus	11
2.2.2	Phoebe	12
2.2.3	Distance and Extinction	13
2.3	Photometry Grids	14
2.3.1	Atmosphere Grids	14
2.3.2	Temperature and Intensity Comparisons	15
2.3.3	Varied Surface Gravities	16
2.4	Test Photometry	17
2.4.1	Preliminary Parameters	17
2.4.2	Limb-Darkening Laws	18
2.4.3	Gravity Darkening	20
2.4.4	Comparison of Light Curves	26
3	LAMOST J112306.9+400736	31
3.1	Previous Study of J1123	32
3.1.1	Reduction of Parameter Spaces	34

3.2	Preparation of Data	35
3.2.1	Reformatting the ZTF Data	36
3.2.2	Conversions	37
3.2.3	Binning	38
3.2.4	TESS Data Preparation	39
3.3	Fitting Optimization	40
3.4	Models	42
3.5	Icarus Model A	47
3.5.1	Comparison with TESS Data	47
3.5.2	Summary	48
3.6	Phoebe Model A	49
3.6.1	Comparison with TESS Data	50
3.6.2	Summary	51
3.7	Icarus Model B	51
3.7.1	Comparison with TESS Data	52
3.7.2	Summary	52
3.8	Phoebe Model B	53
3.8.1	Comparison with TESS Data	53
3.8.2	Summary	54
3.9	Icarus Model C	55
3.9.1	Comparison with ZTF Data	55
3.9.2	Summary	56
3.10	Phoebe Model C	57
3.10.1	Comparison with ZTF Data	57
3.10.2	Summary	59
3.11	Comparison	59
4	Discussion	61
4.1	Physical Accuracy	61
4.2	Versatility	64
4.3	Speed	64
4.4	Overall	65

5	Conclusions	66
A	Parameter Relation Derivations	72
A.1	Mass Estimation Equations	72
A.1.1	Icarus	72
A.1.2	Phoebe	72
A.2		73
B	Corner Plots	74

“This is to certify that the work I am submitting is my own and has not been submitted for another degree, either at University College Cork or elsewhere. All external references and sources are clearly acknowledged and identified within the contents. I have read and understood the regulations of University College Cork concerning plagiarism and intellectual property.”


Sorcha Riain Kennelly


Date

Abstract

Binary star systems known as spiders consist of a millisecond pulsar ablating or accreting from its main sequence companion star. The pulsars in these systems are known for being quite massive and, as such, are useful in the effort to constrain the equation of state for nuclear matter.

Accurate mass measurements of the neutron star in these systems are vital to the field, and are made using bespoke modelling software that simulate the system and generate model light curves. Through fitting these simulated curves to data, accurate estimates for the mass can be made. For years, various codes have been used independently but recently, a question about whether these codes will predict the same systemic parameters was raised.

Here, we present the answer to this question in the context of the two most common spider modelling codes, Phoebe and Icarus, following the first ever comparison between them.

We find that the codes behave quite similarly, as is expected, though there are some differences in how the surface gravity of the companion star is handled, which has some effects on the temperature predictions.

We further explored whether the codes converge to the same conclusions when fitting data by considering the system LAMOST J112306.9+400736 using various constraints and comparing to a previous study by Yi et al. (2022) which used Phoebe. We find that both codes predict similar parameters, though there were some significant differences in the time required to fit a model and in how constrained a model needed to be.

The ultimate choice between codes is personal to each researcher and neither choice will lead to less trustworthy results.

List of Figures

1.1	$P - \dot{P}$ diagram	4
1.2	Proposed equations of state for nuclear matter	6
1.3	Mass estimates for heavy neutron stars	8
2.1	Amplitude variations at different inclinations	10
2.2	Specific Intensity versus Temperature at $\log g = 5$	15
2.3	Specific Intensity versus Temperature at $\log g = 3.5$	16
2.4	Specific Intensity variations with temperature and viewing angle for Icarus(left) and Phoebe (right)	18
2.5	A residual colour plot showing the difference in specific intensity between Icarus and Phoebe at different temperatures and viewing angles.	19
2.6	Icarus 2D surface temperature map	21
2.7	Phoebe 2D surface temperature map	21
2.8	2D surface temperature residual plot	22
2.9	2D log surface gravity residual plot	23
2.10	2D radius residual plot	25
2.11	Test Photometry with identical filling factors	26
2.12	Test Photometry with identical filling factors and gravity darkening disabled	26
2.13	Test Photometry with different filling factors	27
2.14	Matching median temperature and filling factor of 0.6	29
2.15	Matching median temperature and filling factor of 0.95	29
3.1	Previous modelling attempts for J1123	32
3.2	Reformatting and clipping ZTF data	37

3.3	Binned TESS data	39
3.4	Icarus Model A Light Curves	48
3.5	Phoebe Model A Light Curves	50
3.6	Icarus Model B Light Curves	52
3.7	Phoebe Model B Light Curves	54
3.8	Icarus Model C Light Curves	56
3.9	Phoebe Model C Light Curves	58
B.1	Icarus Model A Corner Plot	75
B.2	Icarus Model B Corner Plot	76
B.3	Icarus Model C Corner Plot	77
B.4	Phoebe Model A Corner Plot	78
B.5	Phoebe Model B Corner Plot	79
B.6	Phoebe Model C Corner Plot	80

List of Tables

2.1	Sample star parameters used in testing the binary modelling codes	17
3.1	The Yi et al. (2022) parameters for J1123	33
3.2	J1123 Model Priors	45
3.3	The MCMC parameter values for each of the models.	46

Symbols and Constants

χ^2 : the fit parameter.

P : the orbital period of the system.

$\dot{P}$: the first derivative of the orbital period.

I : the specific intensity.

μ : the viewing angle.

$\log g$: the log of the surface gravity of the companion star.

T_{eff} : the effective temperature.

T_{base} : the base temperature.

T_{irr} : the irradiation temperature.

i : the inclination of the system, given in radians.

q : the mass ratio between the two objects in the binary system, $q = \frac{M_{\text{NS}}}{M_{\text{comp}}}$.

f : the Roche Lobe filling factor of the system.

f_{VA} : the volume-averaged Roche Lobe filling factor of the system.

k_{comp} : the radial velocity amplitude of the companion star in the system.

Ω : the co-rotation factor, which is always 1.

β : the gravity darkening coefficient.

A_{v} : the extinction to the system.

a : the semi-major axis of the binary.

r_{equiv} : the equivalent radius of the companion star.

$abun$: the abundance of heavy elements in a star, relative to the amount in the sun

Z : the metallicity of the star.

Ψ : the gravitational potential.

λ : wavelength.

G : Newton's gravitational constant, $G = 6.67430 \times 10^{-11} \text{m}^3 \text{kg}^{-1} \text{s}^{-1}$.

c : the speed of light, $c = 2.99792458 \times 10^8 \text{ms}^{-1}$

Units

$M_{\odot}$: solar mass, which is the equivalent to the mass of the sun in kilograms - 2×10^{30} kg.

$R_{\odot}$: solar radius, which is the radius of the sun 6.955×10^8 m.

kg : kilogram, the SI unit for mass.

m : metre, the SI unit for length.

s : second, the SI unit for time.

K : Kelvin, the SI unit for temperature.

m/s : metres per second, the SI unit for velocity.

km/s : kilometers per second.

kpc : kiloparsec, one thousand parsecs. A parsec is about 3.086×10^{16} m.

nm : nanometre, 10^{-9} m.

rad : degrees in radians.

$^{\circ}$: degrees.

Acronyms and Abbreviations

dm: distance modulus

DR: data release

EDR: early data release

EoS: equation of state

HJD: heliocentric Julian date

LIGO: Laser Interferometer Gravitational-Wave Observatory

MCMC: Markov Chain Monte Carlo

MSP: millisecond pulsar

NS: neutron star

RL: Roche lobe

SDSS: Sloan Digital Sky Survey

SED: spectral energy distribution

SNR: signal-to-noise

TESS: Transiting Exoplanet Survey Satellite

tMSP: transitional millisecond pulsar

T-O-V: Tolmann-Oppenheimer-Volkoff

ZTF: Zwicky Transient Facility

For Uncle Foxy, who passed before this thesis was fully completed.

He was always so proud of me for following my love of the stars, I wish he could've been here to see this. He would've loved to read this thesis too - and then have had me explain it all in good, simple English to him (or Irish, if I was feeling brave).

Look, Fox, I did it! Yahoo!

Acknowledgements

The most massive thank you to my supervisor, Dr Mark Kennedy, for everything he has done to help me get to this point. He has been helpful, patient and kind throughout his time as my supervisor; he has been an excellent guide and his belief in me has helped me to believe in myself. He has helped me overcome my fear of public speaking, and helped me get involved in Astronomy on Tap and with Cork Astronomy Club. Thank you, Mark, you've been the best supervisor a gal could ask for!

To the Astro lecturers , Dr Michael Tremmel and Prof Bertram Bitsch, thank you both for being excellent supporters of me over the past few years, and for all the wonderful chats at group meetings, journal clubs and astro drinks. An extra thank you to Michael for Astronomy on Tap - Cork, which has become something I care a lot about and hope to continue to support in the future.

Thank you to Prof Paul Callanan for the use of Callisto in 104D and for his support as I have pursued both of my degrees in UCC and in my pursuit of a PhD. Thank you to Danny O'Riordan for the use of his computer in G23, allowing me to run my Phoebe-A model for several months.

Massive thanks to Ben and Lee, my officemates and two of my best friends. You guys have kept me sane, calm and let me use the heaters as much as I needed to keep myself warm. Thanks for playing the games with me every day, and for all the wonderful chats. Thanks for helping me through the harder days, and celebrating with me on the better days. I cannot wait to see how far we all go, and to come home to UCC and share our stories.

Thank you to the Extended Shenanigang, Patryk, Ella, Lilian, Matylda and Raph, for all the laughs in the early days of this degree, and for all the wonderful visits since moving on to pursue further studies abroad or work in industry. Thanks

also to the "wee'uns", our interns and FYP students, Barry, Aisling, Amal, Narcissa and Ethan, who spent time in G23 and came to so many of the astro-nights.

A huge thank you to my best friends, my housemates, my Echo Chamber, Chaia and Ava, for all their help and support over the years and through the best and worst of times. I am so grateful for all the lifts to college from Ava, and all the chats as we walk in or walk home with Chaia, and for the help they both provided with de-stressing and for the constant encouragement.

Also, to Brian, Eibhlín and Aoife, my friends who are a little further away, but who chat with me every day and who have supported me during this project with phone calls, chats and memes: thank you, guys.

Finally, a thank you to my parents, my sister and my brothers, for putting up with my ramblings about the stars and joking about how I will be a forever student. Bonus thanks to my parents for deciding to build their house under a beautiful sky in a beautiful, rural part of the country, kickstarting this whole journey for me. Also, thanks to my dogs for their cuddles every time I come home.

Chapter 1

Introduction

This study focuses on the modelling of binary star systems consisting of a millisecond pulsar and a main sequence companion star. Binary stellar modelling codes, such as Icarus (Breton et al. 2012) and Phoebe (Prša & Zwitter 2005, Prša et al. 2016), have been used to study and model these systems previously (Clark et al. 2020, Kennedy et al. 2022, Yi et al. 2022, Yap et al. 2023) though never directly compared; here, the codes and the results of modelling light curves for these systems containing millisecond pulsars with main sequence companions are directly compared for the first time. This chapter introduces spiders, and the neutron stars within them.

1.1 Neutron Stars

One of the endpoints of stellar evolution is a neutron star. Formed after a star of mass between $10M_{\odot}$ and $25M_{\odot}$ undergoes a supernova explosion (Woosley et al. 2002), a neutron star is a small, dense star supported by neutron degeneracy pressure.

Neutron stars' dense material, nuclear matter, is supported against gravitational collapse by a quantum mechanical force known as neutron degeneracy pressure, derived from the Pauli exclusion principle. The neutrons in the core of a neutron star are tightly compacted which forces some neutrons to occupy higher quantum mechanical states and creates an outward force that supports the structure

of the star against gravity. However, as with white dwarfs, there is a maximum mass that a neutron star can have before the degeneracy pressure is no longer enough to prevent gravitational collapse. For neutron stars, this is the T-O-V limit (Tolmann-Oppenheimer-Volkoff, Oppenheimer & Volkoff 1939), a theoretical maximum mass that a neutron star may have.

Neutron stars were theoretically proposed to exist in 1933 (Baade & Zwicky 1934), following the discovery of the neutron in 1932 (Chadwick 1932). They were observationally detected in 1967 (Hewish et al. 1968).

Since this discovery, numerous neutron stars have been discovered in a variety of different systems, offering astronomers an opportunity to study these ultra-compact objects and try to understand the matter of which they are comprised.

1.2 Spider Pulsars

Pulsars are rapidly rotating neutron stars that emit beamed electromagnetic radiation from their magnetic poles. They are primarily detected at radio wavelengths (Hewish et al. 1968, Arzoumanian et al. 1994).

They rotate rapidly about their spin axis - some complete rotations on the order of milliseconds, and these are known as millisecond pulsars or MSPs (Backer et al. 1982). MSPs can also be known as recycled pulsars as the neutron star is being spun up (or "recycled") by the accretion process (Alpar et al. 1982). Most detected MSPs have been in binary systems known as spiders, wherein the MSP orbits with a main sequence companion star (Manchester et al. 2005). The name is derived from the behaviour of the pulsar in these systems, which are either accreting mass from or ablating their companion, essentially slowly destroying or consuming it. The first spider, dubbed the Black Widow pulsar, was discovered in 1988 (Fruchter et al. 1988).

The region around both stars within which material is gravitationally bound to that star is known as the Roche lobe. The companion star in spiders is typically distorted to match the teardrop shape of this region. Ablation results from the pulsar wind blasting material from this gravitationally bound region, whereas accretion occurs when the companion star overfills this region and material falls towards the pulsar. A measure of how full this Roche lobe is is known as the filling factor,

which is a ratio of the radius of the companion star to the distance from the centre to the L1 Lagrange point, which lies between both stars (Paczynski 1971).

Overtime, it is believed that the spider pulsar will erode and consume its companion completely. Spiders would act as a step in the evolution of isolated pulsars (Ruderman et al. 1989).

Recycled pulsars gain angular momentum through these destructive processes, however they also lose angular momentum. The axis about which the pulsar spins and the magnetic poles are not aligned with one another, which slows rotation through magnetic dipole radiation (Condon & Ransom 2016).

The rotation period, P , and its first derivative, $\dot{P}$ are used to show the evolution of pulsars. The pulsar population, seen in Figure 1.1, highlights that the rapidly rotating pulsars of spider systems are separate to the general neutron star population. Overtime, pulsars slow and move to the right on the diagram. The characteristic age of a pulsar, τ , is then derived from these P and $\dot{P}$ values to be $\tau = \frac{P}{2\dot{P}}$ (Lorimer & Kramer 2004, Jiang et al. 2013). The characteristic ages of these populations is also highlighted in Figure 1.1, taken from Hui & Li (2019), and it indicates that spiders are actually *older* pulsars. They have had more time, as such, to evolve in a binary and develop such rapid spin periods.

In Figure 1.1, there are two separate classes of spider shown, black widows and redbacks. The black widow class of spider binaries contain a companion star with a mass ranging from $0.01M_{\odot}$ to $0.1M_{\odot}$. A second class, known as redbacks, has also been identified (Roberts 2013). They are very similar to black widows, but their companion star is more massive - larger than $0.1M_{\odot}$ - and usually a semi-degenerate main sequence star.

Resultant from the companion star bloating and filling - or partially filling - its Roche lobe, and likely being tidally locked to the pulsar, the companion has a distinct tear-drop shape. This is Roche-lobe distortion. In order to maintain hydrostatic equilibrium on the surface, there is variation in the surface gravity and temperature. At the 'nose', which faces the pulsar, the temperature is lower. Similarly, the back of the companion is also a lower temperature for the same reason. This phenomenon is known as gravity darkening. This temperature variation across the surface of the star gives rise to a variation in the light observed from this system. This creates an optical light curve with two peaks and two troughs - the troughs




Figure 1.1: The $P - \dot{P}$ diagram, where P is the spin period and $\dot{P}$ is the first derivative of the period, plotted using logarithms to show the population distribution of pulsars. The solid black lines denote the characteristic ages, τ , of the pulsars. The dotted lines denote magnetic flux density, and the thick, dashed line marks the region beyond which few to no pulsars are found. The spiders population are denoted by the blue and green circles, while other pulsars are represented by black dots.

corresponding to the colder front and back of the star, and the peaks being the hotter sides (Lucy 1967, von Zeipel 1924).

These light curves are what allow for spider binaries to be modelled, using bespoke spider binary synthesis codes, in order to constrain mass estimates for the stars in the system.

The redback optical light curves feature a characteristic asymmetry, wherein their peaks and troughs do not align perfectly when the light curve is flipped, with variations in their amplitude (Stringer et al. 2021).

These redback systems are particularly interesting as a number of peculiar irregularities have been noticed. Some millisecond redback systems have been observed to transition between x-ray and radio pulsar states (Archibald et al. 2009). There are three canonical tMSPs, and several candidates, all redback systems (Kyer et al. 2025). The cause of the transition is, as of yet, unknown. The best way to confirm a pulsar's status as a tMSP is to observe the transition - as was done for the canonical tMSPs.

Spiders are critical in shedding light on the evolution of pulsars, and more widely neutron stars, and explain how isolated MSPs come to be alone. Due to their nature, these isolated MSPs likely have been spiders in order to explain their existence, while this is not the case for isolated long period pulsars.

Spider masses thus far observed have been on the heavier side of what a pulsar is believed to be able to be. As such, these systems are particularly interesting to study as they may lie right on the believed upper limit of neutron star masses (Linares 2020, Strader et al. 2019, Romani et al. 2022).

1.3 Spiders and the Neutron Star Equation of State

The equation of state for nuclear matter in neutron stars remains elusive. This equation would explain the connection between numerous properties of the neutron star, including a relationship between the mass and radius of the star. This equation of state would provide a maximum possible mass that any neutron star can have before neutron degeneracy pressure would be insufficient to support the stellar structure, resulting in collapse.




Figure 1.2: A plot of the mass versus radius for neutron stars, where the black lines represent a number of proposed equations of state. This image is taken from Paerels et al. (2009)

As is seen above in Figure 1.2, there are numerous proposed equations of state. To constrain which of these is the equation of state, we try to place data points amongst these equations to determine which ones best predict the observed values.

However, the radii of these objects are on the order of 10km, which is incredibly challenging to observe in most cases. Most often, a single data point cannot be accurately placed on a plot, such as that in Figure 1.2, though attempts have previously been made to do so (Ludlam et al. 2022, Vinciguerra et al. 2023), but a horizontal line representing a mass value can be used to eliminate equations that do not predict such a mass. For this reason, accurate mass estimates - especially high and low masses - are of interest as they eliminate particular proposed equations.

High mass neutron stars have been found to occur frequently in spider systems, making them particularly useful in working towards an equation of state, as is seen in Figure 1.3.

Accurate mass measurements are achieved in various ways, but for the purpose of this research, we will delve into fitting models to observed stellar light curves.

1.4 The Mass Gap Question

There has been a gap observed in the masses of compact objects, between neutron stars and black holes, approximately between 3 and 5 $M_{\odot}$ (Özel et al. 2010, Farr et al. 2011, Kreidberg et al. 2012). Is this gap the result of differences in how these objects form or does it instead arise from observational biases? There may be a continuum of compact objects produced and through observational limitations, some may be undetectable from Earth (Kreidberg et al. 2012).

Recently, the lower bound to this mass limit has been challenged. Gravitational wave events have involved compact objects with masses that may lie within this gap (Abbott et al. 2023). One event recently detected by LIGO indicated pre-merger objects of masses within the 2.5 – 4.0 $M_{\odot}$ range (LIGO Collaboration et al. 2024), and a further study by Barr et al. (2024) has discovered a compact object with a mass between 2.09 and 2.71 $M_{\odot}$. These detections challenge both the mass gap’s existence, and push the boundary of the lower limit, should the gap exist.

Spiders, as the hosts of heavy neutron stars, are key to constraining a lower limit for this mass gap, should it exist, through establishing how massive a neutron star can be.

In this thesis, we examine whether the modelling codes which are used to infer the masses of neutron stars in spider binaries are consistent with each other, and how they compare when applied to the same data.

In Chapter 2, the codes used to model the light curves of spider systems are compared to determine if they have similar configurations and investigate whether they will generate models in similar ways.

In Chapter 3, the codes are applied to observational data for a spider to investigate what constraints are required in order to find a best-fit model.

In Chapter 4, the links between the previous chapters will be discussed, and subsequently, conclusions will be drawn on whether or not the binary modelling codes can be used and compared freely.




Figure 1.3: Shown in the top panel are mass estimates for heavy neutron stars. The dashed curves show the two heaviest white dwarf-pulsar binaries. The solid curves show the mass estimates for four black widow pulsars and one red back. The bottom panel shows the normalized likelihood versus maximum neutron star mass for various combinations of these measurements. This image is from Romani et al. (2022)

Chapter 2

Binary Modelling

There are numerous software packages that exist with the purpose of generating synthetic light curves for a given set of binary parameters, intended to be fit to astrophysical data to determine the nature of a system. The choice of code often depends on the science goals of the study, the quality of the data and the accessibility of the code. Codes like ELC (Orosz & Hauschildt 2000) and LCURVE (Copperwheat et al. 2010) have been used for decades to infer information about eclipsing binary star systems. For the duration of this study, the word "binary" is used to refer to a spider binary system, a neutron star with a main sequence companion, unless otherwise stated.

For spider binaries, the primary software choices are Icarus (Breton et al. 2012) and Phoebe (Prša & Zwitter 2005, Prša et al. 2016), which have never been directly compared. Recently, it has been asked whether these codes would produce the same models when presented with the same data and what impact any deviations between these codes would have on established mass measurements (Sen et al. 2024). The system studied in Sen et al. (2024) was initially modelled with Phoebe, but when newer data for the same system was modelled with Icarus, there were notable differences in the predicted binary parameters - such as the predicted inclination, which was 78° in the previous study (Yap et al. 2023), but ranged from 59° - 64° in Sen et al. (2024), having a knock-on effect on the mass estimates. Though some deviation is to be expected as this system exhibits variations in its light curve over long timescales, such large differences in predicted

parameters could indicate a disagreement between the codes.

In the event that this is the case, all measurements derived from these modelling codes are called into question and any previously established trust in these estimates is no longer present.

2.1 Photometry

Changes in the viewing angle of the companion star in a binary system over an orbital period can lead to a periodic variation in the observed flux. This modulation in the observed flux varies across the different filters, resulting from the temperature dependence of the observed colours, and information about the binary system is encoded in these data, such as the inclination of the system, the mass of the stars, the companion star's surface temperature and its Roche lobe filling factor. If the companion star is spherically symmetric or the inclination of the system is low, a small amplitude is observed. An example of this is shown in Figure 2.1, simulated using Icarus.




Figure 2.1: *The variation in the amplitudes of the light curves of spider binaries at different inclinations, shown across five SDSS photometric bands.*

Most spider light curves exhibit ellipsoidal modulation - a characteristic double peak resultant of periodic variations in the brightness of the companion star, which results from the different temperatures that arise on the star's surface and the area

of the surface that is visible. The Roche lobe distorted companion stars in spider systems are assumed to be tidally locked, meaning the same part of the companion star always faces the pulsar. This tidal locking also gives rise to a night and day side of the companion stars. The combination of temperature and geometry effects produce photometric modulations which are perfect for modelling with bespoke synthesis codes.

2.2 Modelling Codes

The python based Icarus and Phoebe codes were compared for this study focusing on the "under-the-hood" behaviour of both.

2.2.1 Icarus

Icarus (Breton et al. 2012) is a python-based binary modelling software intended for modelling binary systems containing a neutron star.

There are several approximations made by Icarus, which improve the speed of model creation but may impact the physical accuracy, such as: approximating the neutron star as a point mass, not including time-dependent and relativistic effects, interpolating between simulated data points, and assuming circular orbits.

Icarus deals exclusively with the companion star when creating its models. For the companion, it calculates the temperature and surface gravity for every viewing angle at every phase.

In order to generate a spider system, Icarus requires several parameters: the mass ratio between the stars ($q = \frac{M_{\text{NS}}}{M_{\text{comp}}}$, the ratio of the stellar masses of the binary), the distance modulus to the system (dm, which is the distance to the system defined on a logarithmic scale based on the magnitude of the star), the co-rotation factor (which is the ratio of the star's spin to its orbital spin frequency and is always set to be $\Omega = 1$), the extinction (A_v , the absorption or scattering of light from the star due to interstellar dust in the Johnson v-band), the Roche lobe filling factor (f), the base surface temperature (T_{base}), the projected radial velocity amplitude (K_{comp} , the projected radial velocity amplitude of the companion star, which is used in the calculation of the masses of the stars in the binary), the orbital period

(P), the inclination (i , which is the angle made by the line of sight of an earth based observer and the orbital axis of the binary), and the exponent for the temperature dependence for gravity darkening (β , which will be discussed further in Subsection 2.2.2). β is set to be 0.25 for stars that have radiative cores, and 0.08 for stars with fully convective cores, in line with Lucy (1967) and von Zeipel (1924).

Icarus can apply a number of heating models to the surface of the companion star in order to explain the observed light curves, including direct irradiation, the ability to add star spots (Clark et al. 2020) and add complex heating effects, such as diffusion (Voisin et al. 2020).

2.2.2 Phoebe

Phoebe is an open source, python-based modelling software that can be used for modelling a variety of binary systems.

Phoebe is the "industry standard" for binary modelling, and can deal with a variety of eclipsing binary systems (Yap et al. 2023, Tobin & Berrington 2024, Cunningham et al. 2020).

Phoebe generally works with normalised flux in all light curves, meaning that some crucial system information is lost - the distance to the system and the extinction. There are some benefits to this approach as working with the distance, as Icarus does, assumes the colour of the companion star matches the colour of a main sequence star, as is assumed in photometry grids. Working with the normalised flux avoids this assumption and also avoids imperfect flux calibrations, which can arise in photometry grids. Icarus does have this same ability to work without assuming colour or distance information, however it is not the intended way for the code to be used.

Phoebe requires most of the same parameters as Icarus to generate a model - such as the inclination, effective surface temperature, mass ratio (though the mass ratio is inverted in Phoebe to be $q = \frac{M_{\text{comp}}}{M_{\text{NS}}}$), orbital period. Phoebe also takes the semi-major axis of the binary orbit (a), the equivalent radius (r_{equiv} , which is the radius the star would have if it were spherical and not distorted) of the companion, and a slightly different version of the gravity darkening exponent.

As previously mentioned, gravity darkening explains the variations in the tem-

perature as the radius of the companion star varies. It is vital in the creation of a model in either code, but both codes do not take the same version of this exponent. The gravity darkening exponent's effect can be parametrised in terms of both L , the luminosity, and T , the temperature, using the Stefan-Boltzmann law as the equivalence:

$$L \propto T^4 \quad (2.1)$$

where Phoebe work in terms of the luminosity, L , and Icarus works in terms of temperature, T , meaning Icarus works with a gravity darkening coefficient that is four times smaller than what Phoebe requires.

Stellar spots and other heating models, such as convective heating or diffusion, can also be applied, as can be done with Icarus.

As a result of accounting for a number of different effects, Phoebe is slower to compute a model and the fitting process is both computationally expensive and time consuming, which will be addressed further in Chapter 4.

2.2.3 Distance and Extinction

Information lost to normalising light curves includes the distance and extinction. As Phoebe works exclusively with normalised data and light curves, it is important to understand what Phoebe may gain in terms of accuracy.

Icarus works with a distance modulus and the extinction in terms of the Johnson V band. Icarus also accounts for calibration issues in the ATLAS9 grids.

It is standard when using Phoebe to work with normalised light curves, which eliminates the calibration issues associated with the photometry grids, though there is information lost in the process (distance and extinction).

The desired results - the mass of the stars in the system - can still be measured, though alternative means of determining the distance and extinction are needed beyond modelling. Through the use of Green's dust maps (Green et al. 2019) and Gaia data releases (Gaia Collaboration et al. 2021), this can be done so this is not a significant issue.

2.3 Photometry Grids

When modelling photometric data, the use of pre-defined atmosphere grids is vital in learning information about distant stars. These grids are the building blocks for models, and this section will explore how Icarus and Phoebe function when working with the same photometry grids, and how this affects parameters such as surface temperature and surface gravity.

2.3.1 Atmosphere Grids

The specific intensity, $I(\mu)$, is the intensity at a given viewing angle ($\mu = \cos(i)$), at a given distance in a given direction. Specific intensity grids consist of a range of three parameters: temperatures (T), surface gravity ($\log(g)$) and viewing angle (μ).

ATLAS9 (Castelli & Kurucz 2003) solves the structure of a stellar atmosphere, which it then uses to predict the specific intensities for a star. A vital aspect of generating models and light curves in both Phoebe and Icarus is providing them with the necessary intensity grids.

Icarus uses ATLAS9 grids as described in Kennedy et al. (2022). Phoebe, on the other hand, required grids to be created. Phoebe has a readily available set of atmosphere tables, dubbed the ck2004 tables, which use grids of specific intensities generated using ATLAS9.

Bandpasses using these grids were generated for both Icarus and Phoebe using their respective in-built functionality for creating band-specific intensities. These were then compared to ensure that these generated bandpasses were similar enough that models would not be significantly different between both codes.

These in-built functions work by multiplying the atmosphere grid with the known bandpass response, and integrating this results in the flux (Bessell & Murphy 2012) This can be seen in the equation below, for a given filter (in this example, the B filter), the flux, f is:

$$f_B = \frac{\int S_B F_\lambda d\lambda}{\int S_B d\lambda} \quad (2.2)$$

Where for a given filter, B , S_B and F_B are the spectrum, F is the flux and S is

the photon response function. λ is the wavelength.

2.3.2 Temperature and Intensity Comparisons

For each SDSS passband filter (u, g, r, i and z), specific intensity versus temperature graphs were generated to ensure the resulting photometry grids were consistent. Once converted to be in terms of specific intensity as a function of frequency, the intensities were plotted over a range of temperatures. The results of this can be seen in Figure 2.2.




Figure 2.2: In the top panels, the red lines represent the specific intensities at a given temperature for the *Icarus* filters, and the blue lines represent *Phoebe*'s specific intensity at a given temperature. The bottom panels show the fractional differences between the intensities at a given temperature, with deviations up to $\approx 5\%$ at low temperatures and $\approx 0\%$ at higher temperatures. The deviation begins below $\approx 4000K$. A $\log g$ of 5 was assumed, and a viewing angle of $\mu = 1$.

Some small deviations are observed in temperatures ranging from $\approx 7000K$ to $\approx 10000K$. This will not be significant in terms of flux or magnitude observed, so it can be concluded that these passbands and the methods by which both codes handle them are similar.

2.3.3 Varied Surface Gravities

The surface gravity is the acceleration due to gravity on the surface of a body. In the modelling softwares, it is normal to work with the log of this value - the $\log(g)$.

The surface gravities 3.5 and 5 were compared to determine if this had any affect on how both softwares aligned. It was found that varied surface gravity, as long as it is the same in both Phoebe and Icarus, does not have an affect. This is seen in Figure 2.3, which shows the codes behave the same regardless of the log gravity value set. There are higher deviations observed at lower temperatures, but they are otherwise consistent.




Figure 2.3: In the top panels, the red lines represent the specific intensities at a given temperature for the Icarus filters, and the blue lines represent Phoebe's specific intensity at a given temperature. The bottom panels show the fractional differences between the intensities, with deviations up to $\approx 10\%$ at low temperatures and $\approx 0\%$ at higher temperatures. The difference is most pronounced at temperatures below 4000K. The log gravity is 3.5 and the viewing angle is still $\mu = 1$.

Though there are some differences in the intensity at lower temperatures, the fractional difference is not significant enough to be detectable in most cases, indicating that both Phoebe and Icarus will produce very similar results from specific intensities, though they may have some small differences.

parameter	unit	value
Mass Ratio ($q = M_2/M_1$)	-	7.8
Inclination (i)	$^\circ$	46.4
Roche Lobe Filling Factor (f)	-	0.999
Radial Velocity Amplitude (k_1)	km/s	282
Distance Modulus (dm)	pc	10.5
Effective Temperature (T_{eff})	K	5800

Table 2.1: *Sample stellar parameters used in the generation of surfaces and light curves in this part of the study. The parameters are based on those of PSR J1023+0028 (Stringer et al. 2021), with a more Roche lobe distorted companion in this case.*

2.4 Test Photometry

A light curve is a plot of flux or magnitude variations over time from which we can derive a number of system parameters by matching the observed curve to one generated using a model.

A useful test for these codes was to generate two models and investigate whether they produce the same results or not.

2.4.1 Preliminary Parameters

The initial parameters chosen here are based on, but not identical to, those of PSR J1023+0038 (hereafter J1023), a transitional millisecond pulsar and redback spider system. The parameters used come from a previous study of the system (Stringer et al. 2021), where it was modelled J1023 using Icarus. J1023 is a well-studied system due to its nature as a tMSP.

Adaptations were made to these parameters in order to facilitate "easy" initial models, without major asymmetries in the light curve - a characteristic of J1023's. Previous studies have found complex heating models necessary to adequately model the system, so removing this feature allowed the most simplistic heating model to be sufficient. In this case, no irradiation was used in order to focus on the distortion of the companion star.

The exact parameters are given in Table 2.1.

2.4.2 Limb-Darkening Laws

Icarus's intensities are taken from the ATLAS9 grids, while Phoebe stores only the intensities at the normal and then calculates the intensities using a logarithmic limb-darkening law, as seen in Equation 2.3 (Klinglesmith & Sobieski 1970), with constants x_λ and y_λ derived from interpolating the intensities with the ck2004 atmosphere tables.

$$I(\mu) = I_0(1 - x_\lambda(1 - \mu) - y_\lambda \mu \ln \mu) \quad (2.3)$$

Here, $I(\mu)$ is the intensity at various viewing angles, μ , with I_0 being the intensity when the viewing angle is $\mu = 1$, which occurs when the observer looks directly at the system, and the constants $x_\lambda = 0.6923$ and $y_\lambda = 0.1628$.

The SDSS passband filters were used to investigate the variation in intensity at different viewing angles and temperatures, with surface gravity and abundance kept constant.

These comparisons are seen in Figures 2.4, where $I(\mu)$ for Phoebe were computed using Equation 2.3.




Figure 2.4: A colour plot indicating the variation in the intensities generated by *Icarus* (left) and *Phoebe* (right) at various values for viewing angle, μ , and temperature, Temp. An abundance of 0 and log surface gravity of 5 was used.

Plotting the specific intensities with respect to frequency, both match very closely with one another. There are some deviations at low viewing angles and

low temperatures, as is evident in Figure 2.5




Figure 2.5: A residual colour plot indicating the variation in the intensities between the intensities generated in *Icarus* and *Phoebe* at various values for viewing angle, μ , and temperature, *Temp*. An abundance of 0 and log surface gravity of 5 was used.

At higher viewing angles, both codes agree, however there are discrepancies at very low temperatures and at low viewing angles. These differences will likely be most pronounced on the surface of the star in places where the temperature and viewing angle is lowest - such as the front of the companion star, facing the pulsar. This will be discussed further in the next section of this report.

During this exploration of the variations in the intensity grids, it was noticed that *Icarus*' documentation did not match with what was being observed in its outputs. On further inspection, the *Icarus* documentation was slightly inaccurate and once accounted for, allowed the output of both codes to be compared. The documentation claimed that *Icarus* was working with dex, however it was actually using the natural logarithm.

This shows that *Phoebe*'s approach of calculating $I(\mu)$ from $I(\mu = 1) = I_0$

using an analytic function could introduce errors between the codes, especially at low μ values, when compared to Icarus's approach of using the computed $I(\mu)$ produced by ATLAS9. That is to say, the analytical function Phoebe uses differs from the true specific intensities of the grids at low viewing angles and high and low temperatures. This causes the analytically calculated specific intensities for any surface elements which lie in these regions to be over- or underestimated in terms of their brightness.

2.4.3 Gravity Darkening

Gravity darkening results from a pressure differential on the surface of a star as it maintains hydrostatic equilibrium. Variations in radius, which are expected for a Roche lobe distorted spider companion, give rise to variations in temperature and surface gravity. This phenomenon leads to the gravity and temperature of the surface being lowest at the distorted front of the star, facing the pulsar (Lucy 1967, von Zeipel 1924).

Through modelling the different temperatures and gravities on the surface of a Phoebe and Icarus star, it became apparent that there were significant deviations. A sample star was used to create these surfaces, with the parameters as given in Table 2.1.

The stellar surface temperature, despite having the same initial conditions, deviate significantly, particularly at the front or nose of the companion star, which is heavily distorted.

These surfaces are plotted as flat, two-dimensional surfaces in Figures 2.6 and 2.7. The darkest, central part of these star surfaces is the elongated nose at the front of the star, closest to the pulsar. The furthest right and left parts of the surface are

the back of the star.




Figure 2.6: A 2D stretched surface for the sample star generated using *Icarus*. The central dark spot is located at the stretched front of the star, and the darker regions at the left and right are the back of the star, not facing the pulsar.




Figure 2.7: A 2D stretched surface for the sample star generated using *Phoebe*. The central dark spot is located at the stretched front of the star, and the darker regions at the left and right are the back of the star, not facing the pulsar.

To further highlight this deviation, a surface residual plot was generated - where the surface temperatures from Phoebe were subtracted from the Icarus surface temperatures. The surface of the Phoebe star required some interpolation as there were fewer points initially created by the code. A `scipy` radial basis function interpolator was used. Following this, a residual plot was generated of the surface temperature. This is seen in Figure 2.8.




Figure 2.8: A residual 2D surface plot showing the difference in temperature between the surfaces created by *Icarus* and *Phoebe*, with a noticeable difference located at the front of the star, facing the pulsar.

It is apparent that the temperature differences are mostly uniform aside from right at the centre of the plot, at the tip of the nose of the distorted companion.

As temperature on the surface varies in order to maintain hydrostatic equilibrium, the other parameters of interest are the log of the surface gravity and the stellar radius.

The surface gravity values also exhibit similar deviations to the temperature, as

is seen in Figure 2.9.




Figure 2.9: A residual 2D surface plot showing the difference in log gravity between the surfaces created by *Icarus* and *Phoebe*, with a noticeable difference located at the front of the star, facing the pulsar, in the same location as the largest temperature deviations.

This difference may arise from the methods by which these log gravity values are calculated.

In *Icarus*, the gravity at a given point (x, y, z) is calculated according to the equation:

$$\log(g_{pole}) = \log_{10} \sqrt{\frac{d\Psi^2}{dx} + \frac{d\Psi^2}{dy} + \frac{d\Psi^2}{dz}} = \log_{10} \frac{d\Psi}{dr} = \log_{10} \nabla \Psi \quad (2.4)$$

where $r^2 = x^2 + y^2 + z^2$ and Ψ is the gravitational potential, and defined to be

$$\Psi = \frac{1}{r_1} + \frac{a}{r_2} + qx + \frac{q+1}{2} \Omega^2 (x^2 + y^2) \quad (2.5)$$

where r_1 is $r^2 = x^2 + y^2 + z^2$, r_2 is $\sqrt{r_1^2 + 1 - 2x}$, a is the semi-major axis of the binary, q is the mass ratio of the stars to one another, Ω is the co-rotation factor, which is always $\Omega = 1$ for spider binaries, and x , y and z are the Cartesian coordinates of a point on the surface of the star. x is defined to be along the direction of motion of the companion star, y is along the line that joins the companion star and the pulsar, and z is perpendicular to both, pointing 'up' and away.

In Phoebe, this calculation is approximately the same.

There is the potential for small differences to arise in how these surface gravities are calculated, which may impact the temperature of the star, most extremely at the front 'nose' of the Roche lobe distorted companion, facing the pulsar.

As was seen in the previous section, the grids used differ most significantly in low temperature areas and at low viewing angles, both of which occur at the nose of the star.

However, there is one final parameter that must be examined: the radius. Despite setting up both stars with identical radii (the maximum possible radius for the star in the binary in Phoebe was multiplied by the Roche lobe filling factor, and then Icarus was given the same filling factor, but volume-averaged), the residuals, as in Figure 2.10 indicate that the stars do not generate completely identical stars.

This fractional difference in the radii is quite small but not insignificant as it contributes to the differences seen in the temperature and the log gravity. The impact this difference would have on the model light curves of this sample system

will be discussed towards the end of this chapter.




Figure 2.10: A residual 2D surface plot showing the difference in radius between the surfaces created by *Icarus* and *Phoebe*, with a noticeable difference located at the front of the star, facing the pulsar.

2.4.4 Comparison of Light Curves

The following light curves were initially obtained for the modified J1023 system, using the parameters in 2.1:




Figure 2.11: The top panel shows simulated light curves across five photometric bands for the sample system created with identical system parameters using *Icarus*, shown in red, and *Phoebe*, shown in blue. The bottom panels show the fractional difference between these model light curves, which is $\approx 2\%$.

The gravity darkening coefficient, β , differs between both codes so to ensure that this is not significantly contributing to the deviations seen in Figure 2.11, gravity darkening was "switched off", and β was set to be 0 in both codes.




Figure 2.12: The top panel shows simulated light curves across five photometric bands for the sample system created with identical system parameters using *Icarus*, shown in red, and *Phoebe*, shown in blue, but with the gravity darkening coefficient, β set to 0. The bottom panels show the fractional difference between these model light curves, which is $\approx 2\%$.

The fractional variation seen in Figure 2.13 is still $\approx 2\%$. This suggests that the gravity darkening is not the cause of this discrepancy. Gravity darkening was set to $\beta_I = 0.08$ and $\beta_P = 0.32$ on the completion of this investigation, "switching on" gravity darkening once more.

Following the results of the previous section of this study, some variations in the Roche lobe filling factor in *Icarus* were made, reducing it from 0.999^3 to 0.92, a fractional difference of $\approx 8\%$. This value was arrived at by choosing various filling factor values in *Icarus* until the light curves matched. An expected deviation in filling factor should be accounted for in photometric studies of systems. This change resulted in an *Icarus* model that matched significantly more closely to the *Phoebe* model.




Figure 2.13: The top panel shows simulated light curves across five photometric bands for the sample system created with identical system parameters using *Icarus*, shown in red, and *Phoebe*, shown in blue. The bottom panels show the fractional difference between these model light curves, which is $\approx 1\%$.

It is noteworthy that the minor differences between the curves in 2.11 and 2.13 are approximately 2% for both filling factors, and current data resolutions are not high enough to detect such a discrepancy. This is not unexpected, as the difference

found in Figure 2.10 showed a fractional difference of $\approx 2\%$ also.

Often, the observed companions are not entirely filling their Roche lobe and so this filling factor parameter is much lower.

The extent of this was quantified by focusing on establishing what filling factors and temperatures for Phoebe produces near-identical surfaces when compared to Icarus.

This highlighted two key discrepancies in the set-up of the binary in both codes - the temperature and the filling factor are *different*. Instead of setting values that were identical, the focus became setting different values that produced identical median values across the surface, and that returned identical filling factors when calculated identically from the radius.

Through varying filling factor values, the radius value of the star (as was used in the generation of surface maps such as those in 2.6 and 2.7), was divided by the binary separation to return a value for the filling factor of the companion star calculated in the same way for both codes.

The aim was to find values that produced near-identical values of this filling factor, evaluate the differences in the resulting fractional difference maps for the radius, and to then match the temperatures. For the temperature, the median value across the surface was established to within 1 K.

Generating surfaces and model light curves for these 'matching' systems showed, crucially, that when these median values match, the differences between the resulting light curves are negligible, amounting to less than a 1% difference.

It is evident from Figures 2.14 and 2.15 that the difference is more pronounced for higher filling factors. It should therefore be expected to find some differences between the temperature values and the filling factors returned by both codes. The filling factor values set in both codes differ less at higher values and more as the filling factor decreases.

Near-identical models *can* therefore be created, they just do not take the same

initial parameters in both codes.




Figure 2.14: The top panel shows simulated light curves across five photometric bands for the sample system created with identical median temperatures and filling factors of 0.6, using *Icarus*, shown in red, and *Phoebe*, shown in blue. The bottom panels show the fractional difference between these model light curves.




Figure 2.15: The top panel shows simulated light curves across five photometric bands for the sample system created with identical median temperatures and filling factors of 0.95, using *Icarus*, shown in red, and *Phoebe*, shown in blue. The bottom panels show the fractional difference between these model light curves.

Summary

Icarus and Phoebe work slightly differently under the proverbial hood, so scientists can expect variations in simulated light curves for some parameters. However the deviations found, amounting to $\approx 2\%$ in most cases, are comparable to the accuracy of most data.

Chapter 3

LAMOST J112306.9+400736

In order to quantify the discrepancies uncovered in Chapter 2, both codes were next used to model real world data. The spider candidate LAMOST J112306.5+400736 has an orbital period of 0.27383544 days (Yi et al. 2022). It is a single lined spectroscopic binary, and believed to consist of an early-type M-dwarf and a massive unseen star, which may be a neutron star.

In this case, data from the potential neutron star LAMOST J112306.5+400736 (hereafter J1123, Yi et al. 2022), which is towards the lower limit of the mass neutron stars are expected to have, was used for several different model configurations in both Phoebe and Icarus. J1123 is an ideal candidate for testing the modelling codes and their consistency for three reasons.

- Dust extinction is minimal (Green et al. (2019)). This avoids observed light being reddened and avoids systematic biases in flux measurements.
- The distance to the system is well-defined through parallax by the Gaia Collaboration et al. (2021).
- The light curves are symmetric, as is seen in Figure 3.1, implying that more simplistic heating models will be sufficient. Typically, it is easier to extract properties from symmetrical light curves.




Figure 3.1: *The light curves for J1123 using a simulated TESS light curve, illustrating the symmetry of the system.*

This combination of minimal extinction, well-defined distance and symmetrical curves suggest that the system should be relatively straight forward to model. The symmetry implies it is unlikely that complex heating models, such as those incorporating stellar spots, will be necessary - though it is worth noting that the previous study (Yi et al. 2022) resorted to using hot and cold spot models. This, in turn, simplifies the models required to explain the observations, while a well-defined distance from the Gaia EDR is helpful for modelling with Icarus and provides a trustworthy, known parameter for the system.

3.1 Previous Study of J1123

J1123 was studied in depth by Yi et al. (2022). The compact object in this binary is proposed to be a very light neutron star, existing at the cusp of where white dwarf

parameter	unit	value
Mass Ratio ($q = M_{\text{NS}}/M_{\text{comp}}$)	-	2.04 ± 0.04
Inclination (i)	$^\circ$	$73^{+1.8}_{-1.3}$
Volume-Averaged Roche Lobe Filling Factor (f_{VA})	-	0.729 ± 0.050
Radial Velocity Amplitude (k_{comp})	kms^{-1}	257 ± 2
Distance (d)	pc	318 ± 4
Effective Temperature (T_{eff})	K	3769^{+8}_{-9}
Pulsar Mass (M_{NS})	$M_{\odot}$	$1.24^{+0.03}_{-0.03}$
Companion Star Mass (M_{comp})	$M_{\odot}$	$0.61^{+0.02}_{-0.02}$

Table 3.1: The parameters found by during the previous study relevant to this study of J1123. The Mass Ratio is defined to be the ratio of the more massive, compact object to its lighter companion star. The Roche lobe filling factor measures, as a fraction, how bloated and distorted the companion star is in terms of its Roche lobe.

stars *stop* being found, near the Chandrasekhar limit. The lack of eclipses in the observed light curves support this, as does the proposed inclination of $73^{+1.8}_{-1.3}^\circ$.

The inclination is high enough that eclipses should be visible in the presence of a white dwarf - unless the white dwarf is very cold ($< 10,000\text{K}$) and very massive. At lower temperatures, the white dwarf would not be luminous enough for us to detect it. However, such a white dwarf would be very old, which is inconsistent with evolutionary modelling of the system. Hence it is most likely a light neutron star, with an estimated mass of $1.24^{+0.03}_{-0.03} M_{\odot}$.

The results of the findings in the previous study were used as initial parameter guesses in the Icarus-Phoebe Comparison models. These parameters are summarised in Table 3.1.

Their values for the effective temperature of the companion star were found using SED (Spectral Energy Distribution) fitting, which entails analysing the wavelength distribution of flux from an object. Their distance value is sourced from Gaia, which uses parallax motion and proper motion to constrain how far away the star is (Gaia Collaboration et al. 2021).

In Yi et al. (2022), TESS data were used. The TESS photometric band is quite broad, ranging between wavelengths of 600nm and 1000nm. This broad band makes it challenging to constrain colour information about a star or system, which may have consequences for assumptions made later by this study.

3.1.1 Reduction of Parameter Spaces

The previous study of J1123 made several assumptions during their pre-modelling steps in order to constrain some parameters, reducing the total parameter space that had to be explored during the optimization process.

In order to constrain the mass and radius of the companion star, isochrones were fit to the absolute magnitude and colour of the system. Isochrones are lines on a Hertzsprung-Russel diagram that pass through stars of the same age. While this technique works in the cases of isolated stars, that evolve on their own along these isochrones, J1123 is a binary system and a single star's evolution is likely not adequate at describing its evolution. This overlooks known stages in a binary's evolution such as the common envelope stage and the mass transfer stage, if it has occurred.

Despite these constraints, their model appears to be under-fitting the data. Under-fitting indicates that their model does not sufficiently describe the observed light curves, and the residual plots, as seen in Yi et al. (2022) Supplementary Figure 7A, have a clear structure arising at the same locations the model fails to fully account for the system's curvature.

To solve the under-fitting, where the model does not adequately describe the data, additional hot and cold spot models were considered, as is seen in Yi et al. (2022) Supplementary Figure 7B and C, but as will be seen from the numerous models fit to the ZTF data for J1123, a hot or cold spot is *not* necessary to obtain a good fit for the data and implied that the parameters obtained for the TESS data are reflective of an overly constrained model.

It is evident in Yi et al. (2022) Supplementary Figure 7, the model is noticeably under-fitting at the peaks and troughs without the addition of hot (as in B) or cold (as in C) spots.

Hot and cold spot models have often been used to describe the observed curves in spider systems, such as in Stringer et al. (2021), Clark et al. (2020) and Sen et al. (2024), however there is no direct evidence that hot or cold spots exist on the surfaces of the companion stars beyond these models, where their size, temperature and location are chosen such that the best possible fit is applied to the data. Though it is possible they *do* exist, the models are idealistic in the way in which spots are

applied.

Hot and cold spots are frequently used in redback systems where asymmetries occur in the light curves, which is not the case here. Asymmetries arise when the peaks or troughs in the light curve have different amplitudes, and may indicate the heating on the surface of the distorted companion star is not symmetrical. Furthermore, the locations of the spots in the models presented by Yi et al. (2022) are at the pointed front of the star, and the back. This implies it may be a gravity darkening issue rather than an actual need for stellar spots, as the front (or nose) of the companion is subject to the strongest gravity darkening.

The constraints imposed by the SED and isochrone fitting contributed significantly to this issue, as these are not sufficient for constraining information about spiders.

Using both codes to model the system will determine whether the necessity for spots is code-specific issue. Utilising TESS data may also lead to some loss of useful information for constraining the temperature of the companion star.

3.2 Preparation of Data

Data for the system J1123 have been taken by several sky surveys. For this study, data from the Zwicky Transient Facility DR23 (Bellm et al. 2019, Masci et al. 2019) and Transiting Exoplanet Survey Satellite (Ricker et al. 2015) were used.

The Zwicky Transient Facility (ZTF) is a wide-field sky survey which had its first light in 2017 and has been taking data since. The J1123 data were taken by the ZTF instrument over the course of four years, starting on March 25 2018 and finishing on March 3 2022, with readings taken approximately every three days.

There are three ZTF photometric bands - the g, r and i bands. The g-band has wavelengths roughly ranging from 410nm to 550nm, the r-band ranges from 590nm to 700nm and the i-band ranges from 720nm to 880nm, similar to the SDSS filters used in Chapter 2.

The Transiting Exoplanet Survey Satellite (TESS) was intended to search for exoplanets but incidentally studies a variety of interesting objects as it surveys the sky. TESS data is stored in units of counts per second, so it is common to normalise these data, as is the case for the J1123 data. The TESS data were taken

over a period of several weeks. The first reading was taken on February 19 2020, and the final was taken March 17 2020, with readings taken approximately every 30 minutes.

TESS data is taken across a very broad photometric band, ranging from wavelengths of 600nm to 1000nm, whereas the ZTF bands are much narrower. A consequence of this single, broad band is that it is not possible to constrain colour information about the system being observed, while the ZTF bands provide a more comprehensive understanding of the colour of a star.

The relevant ZTF data and TESS data then underwent some reformatting and clipping to facilitate modelling.

3.2.1 Reformatting the ZTF Data

The phase data for both systems are continuous and stored in *.fits* files, so they were folded in both instances to correspond to a single rotation - known as phase folding. The ephemeris used for this is given in Equation 3.1, taken from Yi et al. (2022). Further to this, the data were arranged chronologically by phase.

$$T(\phi = 0) = 2453734.90932(31)HJD + 0.27383544(18) \times N, \quad (3.1)$$

where *HJD* is the heliocentric Julian date, and *N* is the number of cycles, which is multiplied by the orbital period. The time, T_0 , located at phase $\phi = 0$, corresponds to the superior conjunction of the companion.

The three different photometric bands, g-, r- and i-, were separated next, using the approximate range of magnitudes for each band. There were some additional noisy points, scattered far from the central curve for each band. To establish the significance of these points and ensure the best data points were used for the modelling process with Icarus and Phoebe, a more robust method was necessary.

The astropy (Astropy Collaboration et al. 2013, Astropy Collaboration et al. 2018, Astropy Collaboration et al. 2022) module has an inbuilt modelling function, which was used to fit a fourier series to each band. The residuals were then used to

clip noisy points. Any residual point more than 2σ from the mean was removed.




Figure 3.2: *The light curves of the ZTF data for the system in three photometric bands, with a Fourier model fit to the data used for sigma clipping. The black points indicate the points that were removed following sigma clipping. The y-axis represents the magnitude of the system, and the x-axis is the orbital phase, calculated using Equation 3.1.*

From this point, the data are ready to be modelled using either Icarus or Phoebe.

3.2.2 Conversions

The data obtained from the ZTF database were in magnitudes, so this had to be converted to flux and normalised for use with Phoebe. The relationship between these is:

$$m - M_r = -2.5 \log_{10} \frac{F}{F_r} \quad (3.2)$$

where m and M_r are magnitudes, F and F_r are fluxes and the r indicates refer-

ence values. Equation 3.2 can be rewritten in terms of AB magnitudes.

$$m = -2.5 \log_{10} F - 48.60 \quad (3.3)$$

where $M_r = -48.60$ is the zero point AB magnitude in cgs units. Rearranging this equation gives the conversion used for this study:

$$F = 10^{(\frac{m+48.60}{-2.5})} \quad (3.4)$$

The error associated with each flux datum was calculated also, using the propagation of errors formula:

$$\sigma_F^2 = \left(\frac{\partial F}{\partial m}\right)^2 \sigma_m^2 \quad (3.5)$$

where F is the flux, m is the magnitude, and σ is the error associated with that parameter. Note that the assumption that there was no error in the zero-point magnitude was made. $\frac{\partial F}{\partial m}$ is defined as follows:

$$\frac{\partial F}{\partial m} = (\log 10) \left(10^{\frac{m+48.60}{-2.5}}\right) \left(\frac{1}{-2.5}\right) \quad (3.6)$$

To normalise the data, the flux was then divided by the median flux. For the error associated with each data point, the median flux was also used.

3.2.3 Binning

Data binning is a process by which large data sets, as with this study, are broken into discrete intervals or bins. All data points within the bin are then combined into a median value. This has the benefit of both reducing the number of data points and of reducing the error associated with a given data point.

The ZTF data were not binned during the fitting process for several reasons. Firstly, binning data can create biases and lead to subtle features of the light curves being lost when potentially significant features are combined into a binned point. Additionally, working with the full, raw dataset preserves any potentially correlated noise which is lost when using binned data.

Crucially, when using an MCMC algorithm for the fitting procedure, main-

taining as much information as possible for the likelihood function helps to best constrain the best-fitting parameters.

However, once the modelling process was completed, the accuracy of the fits was made clear using a binned light curve to highlight how well the model matches these data. The data were divided into forty phase bins, and the median flux in this phase bin was used as a data point.

3.2.4 TESS Data Preparation

The TESS data were also phase folded and normalised, as with the ZTF data. The data did not need to be converted from magnitudes to flux as they were already in flux units.

In the previous study, the binned TESS data were used for modelling, but in this study, the unbinned light curves were used for the fitting. Following this, the data were binned for TESS into forty phase bins of equal size to ensure the model adequately described the curve seen.




Figure 3.3: Shown here are binned TESS data, where the grey points represent the unbinned light curve and the blue points form the binned light curve. There are forty phase bins used for the binned curve. There are error bars on both the binned and unbinned data points, but the errors on the binned points are very small.

In binned data, such as those in Figure 3.3, the errors associated with the data points are smaller than the original errors in the unbinned data. This results from the combined points factoring out noise and other fluctuations.

3.3 Fitting Optimization

To ensure that the most probable model which best describes the data is found, an algorithm which cycles through various parameter combinations is used. These return a "goodness-of-fit" parameter, in this case a χ^2 value, which indicates how closely the model and data align.

Bayesian statistics are employed, allowing for pre-existing biases to be considered as the probabilities are evaluated in the determination of the best-fit model. Baye's theorem evaluates the probability of an event *given* another event. In Equation 3.7 the probability of an event, A, occurring given that event B has occurred is shown.

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (3.7)$$

$P(A|B)$ is the posterior probability, $P(B|A)$ is the likelihood, $P(A)$ is the prior probability, and $P(B)$ is the probability that the evidence is true.

Here, a Markov Chain Monte Carlo (MCMC) algorithm is used, which calculates the log posterior probability by adding the log likelihood and the log priors, implementing this Bayesian theory. The probability in 3.7 can be rewritten as:

$$\text{Posterior} = \frac{\text{Likelihood} \times \text{Prior}}{\text{Evidence}} \quad (3.8)$$

It is preferable to work in log space. With L being the likelihood, P being the prior and E being the evidence, the posterior in Equation 3.8 can be written as:

$$\ln \left(\frac{L \times P}{E} \right) = \ln(L) + \ln(P) - \ln(E) \quad (3.9)$$

The evidence can be challenging to estimate but as $\ln(E)$ is constant, it is ignored in minimisation problems such as this. As is shown below, $\ln(L) \propto \chi^2$. Setting $\ln(P)$ to be 0 when the parameter values fall within the specified interval will help to maximise $\ln(\frac{L \times P}{E})$, which will give the best value.

$$\ln(\text{Posterior}) = \ln(\text{Likelihood}) - \ln(\text{Prior}) \quad (3.10)$$

The log likelihood utilises the χ^2 value for the model. This values is calculated for each individual model created, and is derived from Gaussian distributions.

The equation for a Gaussian probability distribution is given by:

$$P(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp - \frac{1}{2} \frac{(x - \mu)^2}{\sigma^2} \quad (3.11)$$

where $P(x)$ is the probability, μ is the expectation value, σ is the standard deviation and x is the measured variable.

For some model that predicts a value $x = x_m$, the probability, $P(x)$, that the model is correct, given some observation $x \propto \sigma$, where σ is the error, can be written as:

$$P(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp - \frac{1}{2} \frac{(x - x_m)^2}{\sigma^2} \quad (3.12)$$

The total probability given n observations is given by:

$$P(x_n) = \prod_n P(x) \quad (3.13)$$

The goal is to then find a model that maximises this probability. Rather than working in linear space, it is easier to work with the log probabilities. This is the log likelihood, as mentioned previously.

$$\ln(P(x)) = \sum_n \ln\left(\frac{1}{\sigma\sqrt{2\pi}}\right) + \sum_n -\frac{1}{2} \frac{(x - x_m)^2}{\sigma^2\sqrt{2\pi}} \quad (3.14)$$

For a fixed σ , that means that the relevant term is just:

$$\chi^2 = \sum_n \left(\frac{x - x_m}{\sigma}\right)^2 \quad (3.15)$$

where x is the value of the data at a given point, x_m is the value of the model at this same point, σ is the error associated with the data point, and n is the total number of points. Equation 3.15 assumes a Gaussian distribution for the data.

The reduced χ^2 is the total χ^2 divided by the total number of degrees of freedom, which is the number of datapoints minus the number of parameters being fit. By minimising the χ^2 value, in conjunction with maximising the prior probability, the probability that the model predictions are correct is maximised.

For Icarus, the *emcee* (Hogg et al. 2010) package was utilised in order to carry out the MCMC fitting. Following the steps as detailed in Foreman-Mackey et al. (2013), Icarus’s model could be fit to data relatively quickly.

With *Phoebe*, inbuilt MCMC functionality exists that also utilises the *emcee* package. This in-built MCMC process was time-consuming. As *Phoebe* does not interpolate model data points like Icarus, computing the model is much slower. When cycling through a large parameter space, for multiple parameters, this process would not be completable within the time constraints of this study.

It was necessary to adapt the same *emcee* process as Icarus for *Phoebe* instead of its in-built MCMC process, while also incorporating interpolation when generating the model to speed up the process. This involved using a reduced time-series of fifty points, which was then expanded using spline interpolation, which essentially generated a smooth curve through the reduced number of data points, to calculate the χ^2 value for a given data set.

3.4 Models

Six different models were tested to determine how the codes compare, what additional constraints are needed when using a code, and to indicate whether one code would be a better choice than the other. Three Icarus models and three *Phoebe* models were generated. In line with the previous section of this study, ATLAS9 atmosphere grids were used. It is unclear, based on the information provided in the previous study, what atmospheres were used.

To robustly test both codes, it was necessary to compare how they performed across various circumstances - such as when additional fitting procedures have

been employed (as was the case with Yi et al. (2022)) and when the model are given minimal amounts of information.

For each code, there is an "unconstrained" run, where as many unknown parameters as possible are fit to the ZTF data for J1123, hereafter referred to as the (A) models. These models represent both situations where not much additional information is known about the system and also useful test models for determining whether the code requires heavy constraints in order to come to the same conclusion as was found in the previous study of J1123. The details of the exact priors used for these models are given in Table 3.2. Only the ZTF data was used for the optimisation of these models as using only ZTF data allowed Icarus to work with raw data (rather than normalised, as is the default for TESS data). This would provide a good test of how well the normalised and raw data fitting compared between both codes.

There is then also a "Yi-restrictions" model for both codes, fitting the same parameters using approximately the same priors as what was done in the Yi et al. (2022) study of J1123. Again, this uses the ZTF data for J1123. These are hereafter referred to as the (B) models, and the priors used to generate them is given in Table 3.2.

The previous study used SED fitting to constrain its temperature, volume-averaged radius, distance, bolometric luminosity and surface gravity, and isochrones to constrain the mass of the companion star to be $0.61 \pm 0.02 M_{\odot}$. These estimates on stellar radius, distance and temperature were used, alongside the companion mass, when setting up the binary models with few, tight constraints on parameters.

Finally, for both Icarus and Phoebe, the TESS data for J1123 were fitted using the "Yi-restrictions" to ensure that the previous results were reproducible. This is known as the (C) model for both codes. The only difference between the data fit by the C-model and the data in the previous study is that the data in the C-model are not binned. There were two reasons for this. Firstly, there may have been some structures in the data that would be lost if it were binned, as noted previously, and secondly, the adapted MCMC algorithm employed for the other Phoebe models interpolates fifty modelled data points to calculate the accuracy of the fit for a given set of parameters, while the number of bins used in the previous study was forty which would mean reducing the number of data points. To maintain

consistency between all the Phoebe models for J1123, this same MCMC algorithm was employed with fifty data points and so binned data, matching that used in the previous study, was not used. Only TESS data was used in this fitting process, not ZTF alongside it, so as to ensure an accurate one-to-one comparison between the previous study and the C-models.

The details of the parameters fit by each model and the types of prior distributions that they used are summarised in Table 3.2.

The semi-major axis is linked to the radial velocity amplitude according to Equation 3.16, and the filling factor is linked to the equivalent radius. The filling factor required by Icarus, which is the radius of the star at the front/nose divided by the distance from the centre of the companion to the L1 Lagrange point, is slightly different to that taken by Phoebe, which is in terms of the maximum equivalent radius the companion star may have. The relationship between the two filling factors is $f_I = f_P^3$ approximately.

Furthermore, Phoebe takes the equivalent radius instead of the filling factor so in order to fit for the filling factor, the maximum radius of the companion had to be calculated and the filling factor could be incorporated this way. In addition to that, the radial velocity amplitude was fit using its relationship with the semi-major axis. This relationship is written as in Equation 3.16.

$$a = \frac{Pk_{comp}(1+q)}{(2\pi)\sin i} \quad (3.16)$$

where a is the semi major axis in km (which must be converted to solar radii for use in Phoebe), k_{comp} is the radial velocity amplitude of the companion star, P is the period of the binary, $q = \frac{M_{NS}}{M_{comp}}$ is the mass ratio, and i is the inclination in radians. Note, however, that Icarus and Phoebe use mass ratios that are the inverse of one another, so the mass ratio constraints applied differ slightly in 3.2.

The results for each model are summarised in Table 3.3 and are visually presented as corner plots in Appendix B.

parameter	Icarus-A	Icarus-B	Icarus-C	Phoebe-A	Phoebe-B	Phoebe-C
Mass Ratio	$\mathcal{G}(1, 20)$	$\mathcal{U}(1, 20)$	$\mathcal{U}(1, 20)$	$\mathcal{U}(0.05, 1)$	$\mathcal{U}(0.05, 1)$	$\mathcal{U}(0.05, 1)$
Inclination	$\sin(0, 90)$	$\mathcal{U}(0, 90)$	$\mathcal{U}(0, 90)$	$\sin(0, 90)$	$\mathcal{U}(0, 90)$	$\mathcal{U}(0, 90)$
Radial Velocity Amplitude	$\mathcal{G}(257, 73)$	-	-	$\mathcal{G}(257, 73)$	-	-
Volume-Averaged filling factor	$\mathcal{U}(0.001, 0.999)$	-	-	$\mathcal{U}(0.001, 0.999)$	-	-
Effective Temperature	$\mathcal{U}(3500, 15000)$	-	-	$\mathcal{U}(3500, 15000)$	-	-
Distance Modulus	$\mathcal{U}(7, 8)$	-	-	-	-	-

Table 3.2: The $\mathcal{U}(x_1, x_2)$ represents a uniform distribution, $\sin(x_1, x_2)$ represents a sinusoidal distribution, $\mathcal{G}(x_1, \sigma_1)$ represents a Gaussian distribution. The priors for the B and C models are taken from Yi et al. (2022).

parameter	Icarus-A	Icarus-B	Icarus-C	Phoebe-A	Phoebe-B	Phoebe-C
q	2.2 ± 0.2	$1.92^{+0.01}_{-0.01}$	$1.62^{+0.03}_{-0.03}$	$3.36^{+0.31}_{-0.04}$	2.26 ± 0.01	2.04 ± 0.02
i	84^{+4}_{-8}	67.4 ± 0.3	88^{+1}_{-2}	85^{+3}_{-4}	87^{+1}_{-2}	$72.9^{+1.6}_{-0.5}$
k_{comp}	257^{+2}_{-2}	-	-	257^{+3}_{-3}	-	-
f_{VA}	$0.71^{+0.02}_{-0.01}$	-	-	$0.83^{+0.16}_{-0.03}$	-	-
T_{eff}	3899^{+38}_{-26}	-	-	7642^{+179}_{-334}	-	-
dm	$7.5^{+0.3}_{-0.2}$	-	-	-	-	-
M_{NS}	$1.05^{+0.15}_{-0.08}$	1.414 ± 0.006	1.26 ± 0.02	$0.82^{+0.08}_{-0.03}$	$1.003^{+0.005}_{-0.003}$	$1.23^{+0.02}_{-0.03}$
M_{comp}	$0.478^{+0.129}_{-0.079}$	$0.736^{+0.001}_{-0.001}$	$0.778^{+0.027}_{-0.027}$	$0.245^{+0.074}_{-0.009}$	$0.443^{+0.002}_{-0.003}$	$0.599^{+0.031}_{-0.023}$
$\chi^2_{red}(g)$	3.91	4.14	3.91	3.87	3.90	4.34
$\chi^2_{red}(r)$	4.18	4.24	4.32	4.18	4.19	5.15
$\chi^2_{red}(i)$	2.32	2.71	2.30	2.29	2.22	2.53
$\chi^2_{red}(TESS)$	0.93	1.37	0.90	2.16	1.76	1.21
χ^2_{red}	11.34	12.47	11.44	12.52	12.07	13.23

Table 3.3: The results from each model are shown here, with the masses of both stars in the binary calculated using the results. The χ^2_{red} for each band, and total χ^2_{red} for each model is included also.

3.5 Icarus Model A: Unconstrained Fit

The first Icarus model has minimal constraints on the values being fitted, with wide priors and six different parameters that it fits for - in addition to calculating a mass for both stars in the system.

The parameter values that the Icarus code was given as initial conditions are summarised in Table 3.2. The priors were mostly flat (equal probability within a given range of values) though the inclination values were weighted with a sine prior as it is more likely that systems with higher inclinations will be detectable (an observation bias), and the radial velocity and mass ratio were binomially distributed.

The Icarus model was run until convergence, and successfully converged (meaning all the walkers in the MCMC were agreeing on the best fit values) after ≈ 7 million steps. The reduced χ^2 value was 11.34. The χ^2_{red} for each band is included in Table 3.3.

The simple, direct heating model used manages to explain the system well, avoiding the need for hot or cold spots. A direct heating model assumes the surface of the companion star is heated by the pulsar directly and that no further heat redistribution, such as convection or diffusion, takes place

The results of the fitting algorithm are summarised in 3.3. The inclination is higher than what had been determined in the previous study, but the mass estimates are consistent. The volume averaged Roche lobe filling factor, f_{VA} , agrees with the previous study, as does the radial velocity, effective temperature and the mass ratio.

3.5.1 Comparison with TESS Data

To test the accuracy of this model, the TESS data used in the previous study were added to the Icarus model *following* the fitting procedure. This essentially allowed the parameters found to be tested against the TESS data, and check that they can adequately describe and model this other light band.

Through analysing the fit to these data and the ZTF data simultaneously, it is clear that the parameters from the MCMC produce a model that does adequately explain the light curves observed.

The reduced χ^2 for the TESS band is below 1, which may imply over-fitting. This is not unexpected, as the Icarus model was generated using distance information which is not needed for normalised data. It is also a possibility that using normalised data may have resulted in some error underestimation, leading to this low χ^2 value.

When the parameter results from the previous study were used in Icarus, the model was also found to be over-fitting the data.




Figure 3.4: The top panels show the light curves across the ZTF g-band, r-band and i-band, and the TESS band with the Icarus-A models fitted to the data. The bottom panels show the residuals for these bands calculated from the binned data.

The residuals for the Icarus-A model seen in Figure 3.4 are largely unstructured in the ZTF bands, though there is some structure evident in the TESS band and potentially also in the i-band.

3.5.2 Summary

The Icarus-A model returns results consistent with the previous study, though the errors on the parameters are bigger, without the need for isochrone and SED fitting to overly constrain various parameters, and didn't require data to be normalised,

preserving distance and extinction information.

The final results align relatively well with the findings of the previous study, aside from the larger inclination, though they are within two standard deviations from one another. The errorbars are larger on the values generated by this Icarus model, though this is to be expected as the MCMC algorithm explored wide parameter spaces.

3.6 Phoebe Model A: Unconstrained Fit

The first Phoebe model has minimal constraints, as with the first Icarus model. Again, the emcee fitting process has wide priors and five different parameters that it fits for - in addition to calculating a mass for both stars in the system using the parameters being varied.

The parameter values that the Phoebe code was given as initial conditions are summarised in 3.2 Most of the parameter priors were flat (equal probability within a given range of values) though the inclination values were weighted with a sine curve, and the radial velocity and mass ratio had a Gaussian distribution applied. This was the same in the Icarus-A model.

The Phoebe-A model was run for as long as possible but had not yet converged after seven hundred thousand steps and several months of running. The discrepancy in the number of steps arises from the difference in the length of time taken to generate a model in Icarus and Phoebe - Phoebe is much slower. The reduced χ^2 value was 12.47, which is higher than for Icarus-A. The χ^2_{red} for each band can be found in Table 3.3.

The parameters found by the MCMC are summarised in Table 3.3. The high inclination was consistent with what Icarus-A predicted, and the companion star's velocity was consistent with the previous study and the Icarus model. The results diverge for the other parameters, indicating a very low compact object mass, a significantly hotter surface on the companion star and a higher mass ratio. This temperature is not consistent with the spectroscopic temperature found in Yi et al. (2022), which arises due to the lack of colour information available when working with normalised light curves.

The MCMC did not converge, and appears to have become stuck in a local minimum for the χ^2 value.

The fits, though not perfect, do not suggest a need for star spots, as with the Icarus-A results.

3.6.1 Comparison with TESS Data

The model's ability to explain both ZTF and TESS observational data is important in ascertaining whether the system parameters are accurate and so, the TESS data were modelled using the results from fitting the ZTF data. In the case of this model, the data were already normalised so no additional steps had to be taken in this regard to incorporate the TESS data.




Figure 3.5: The top panels show the light curves across the ZTF g-band, r-band and i-band, and the TESS band with the *Phoebe-A* models fitted to the data. The bottom panels show the residuals for these bands.

There is obvious structure in the residuals, particularly in the TESS band, showing that this model is inadequate.

The χ^2_{red} value for the TESS band is 2.16, indicating the fit is poorer than the other five models. This is clear in the residuals as the structure is quite pronounced.

3.6.2 Summary

The Phoebe-A model results are not consistent with the previous study, with larger errors. The model, despite running since August 2024, was time consuming to run a fitting algorithm for, and failed to converge after 700000 steps.

These results indicate Phoebe likely requires some additional constraints, as was needed in the previous study, to reduce the parameter spaces being explored during the fitting procedure. This unconstrained run was not successful in getting the model to converge, and most of the parameters returned by the optimisation vary significantly from the other five models and also the previous study.

3.7 Icarus Model B: Yi-Constrained Fit

This Icarus model adopts a more constrained approach to fitting, using the flat priors detailed in the previous study of the system and the smaller, tighter prior ranges.

The previous study only ran for 10,000 steps, so this procedure was followed when generating the Icarus-B model. This was insufficient for the model to converge, even with more restricted parameter spaces.

In the previous study, the companion star's mass was fixed at $0.61M_{\odot}$ and a radial velocity curve was fit simultaneously with the light curve. With Icarus unable to specify a companion star mass, and in the absence of a radial velocity curve for the system, the radial velocity was fixed according to the results from Yi et al. (2022) and the companion star mass was calculated from the model's fitted parameters. The inclination was not weighted with a sine curve, so the B- and C-models all use a uniform weighting for this value.

The inclination differed from that found in the previous study, but the other parameters were more in line with expectations.

The low error bars are reflective of the MCMC not having sufficient time to explore the parameter space and converge.

The χ^2_{red} for this model was 12.47 which is slightly better than the Phoebe-A model and not as good as Icarus-A.

3.7.1 Comparison with TESS Data

Again, TESS data were modelled using the results of the fit generated for the ZTF data, which can be seen in Figure 3.6.




Figure 3.6: The top panels show the light curves across the ZTF g-band, r-band and i-band, and the TESS band with the *Icarus-B* models fitted to the data. The bottom panels show the residuals for these bands.

The χ^2 value for this TESS-band model was 1.37, indicating the fit here is good. However there is some structure evident in the ZTF bands' residuals and a pronounced structure is in the residual for the TESS band, despite the "good" fit.

3.7.2 Summary

The *Icarus-B* model results are not consistent with the previous study, but they are quite close to the same values. The errors for the B-model are quite small, a consequence of the short run in a tight parameter space.

These results indicate that 10,000 steps may be insufficient for *Icarus* to fit these data, though this short run was undertaken for comparisons to the previous study. If it had run longer, the model would likely have converged and also

had more time to sufficiently explore the parameter spaces available which may have returned parameter values more akin to the other models (except, of course, Phoebe-A). This same 10,000 step run was used for all subsequent models, and while allowing them to run until convergence may have produced the best possible results, this shorter run shows how identically set-up models progress after a set amount of time, and facilitates one-to-one comparisons with the previous study.

3.8 Phoebe Model B: Yi-Constrained Fit

This model is the closest to that run in the previous study while still using the ZTF data for J1123. The parameters and the priors are detailed in Tables 3.2 and 3.3.

The model was run for 10,000 steps and failed to converge. As previously mentioned, the past study of this system fixed the companion star mass, but this was not the case for this model. Due to the absence of radial velocity curves, the radial velocity was fixed according to the priors in Yi et al. (2022) and the companion mass was calculated using the results of the fitting algorithm.

This model is notable as it is a replication of the Yi et al. (2022) models but with ZTF data instead of TESS data.

The χ^2_{red} for this model is 12.07, which is better than the Icarus-B and Phoebe-A models, but not as good as the Icarus-A model.

For the most part, the results differ from the previous study's results, with a larger inclination that agrees with the Icarus-A model, a higher mass ratio agreeing with the A model also, and both masses lie within the errors of the Icarus-A results.

This suggests that the ZTF and TESS data lead to slightly different predictions for the inclination.

3.8.1 Comparison with TESS Data

As before, the model's ability to explain both ZTF and TESS data is important and so, the TESS data were modelled using the parameter results for Phoebe-B from

fitting the ZTF data. The results of this can be seen in Figure 3.7.




Figure 3.7: The top panels show the light curves across the ZTF g-band, r-band and i-band, and the TESS band with the *Phoebe-B* models fitted to the data. The bottom panels show the residuals for these bands.

Again, there are some obvious periodic structures in the residual plots, particularly in the residuals for the TESS data, suggesting that the values for this model do not sufficiently describe the system. The χ^2_{red} for the TESS band is 1.76, the second highest of these values, further emphasizing that this fit is not adequate.

3.8.2 Summary

The *Phoebe-B* model results are not consistent with the previous study but they are consistent with *Icarus-A*. The current parameter combination is not adequate at describing the system, leading to structures in the residual light curve plots.

The errors associated with the parameter values are small, but the MCMC only ran for 10000 steps in a reduced parameter space and did not converge. If the MCMC had run for longer and with less constrained parameter spaces, the results might be more trustworthy.

These results emphasize that *Phoebe* requires some additional constraints for

the fitting process, as was needed in the previous study and indicated by the Phoebe-A results, to reduce the parameter spaces being explored during the fitting procedure, but they also suggest that over-constraining these parameters will not eliminate structures from the residuals. Even when the values have small errors, they are clearly not sufficient so the model likely needs looser constraints.

3.9 Icarus Model C - Constrained Fit with TESS data

This model is a replication of the Yi et al. (2022) study using Icarus instead of Phoebe. For this model, the TESS data were used in the MCMC fitting procedure instead of the ZTF data.

As for the B-models, this model was tightly constrained with only the mass ratio and the inclination being varied during the MCMC and the masses of both stars being derived from these values. This model also only ran for 10,000 steps and failed to converge in this short time.

The priors for this model are given in Table 3.2 and the results are in Table 3.3. The value for the mass ratio does not agree with the previous study, nor does the inclination, but the inclination is consistent with Icarus-A and Phoebe-B. The neutron star mass agrees with the previous study, but the companion mass does not.

This model has a χ^2_{red} value of 11.44, which is better than all other models except Icarus-A. The TESS-band χ^2_{red} for this model is less than 1, indicating that there may be some over-fitting happening for this band.

3.9.1 Comparison with ZTF Data

As before, the best test of these parameter results is to see how the models generated using them fit to all four bands. The ZTF data were modelled using the results

of Icarus-C, and this can be seen in Figure 3.8.




Figure 3.8: The top panels show the light curves across the ZTF g-band, r-band and i-band, and the TESS band with the *Icarus-C* models fitted to the data. The bottom panels show the residuals for these bands.

There are no prominent structures in the light curve residuals. The χ^2 values for Icarus-C suggest that the fit is good but still not as good as Icarus-A's values.

However the model had not adequately finished the fitting procedure as the best fitting model should be the TESS data model.

3.9.2 Summary

The Icarus-C model is an attempt to replicate the results from Yi et al. (2022) using Icarus instead of Phoebe. The model failed to converge in 10,000 steps and the light curve residuals indicate that the parameters are not adequate for generating a model that accurately describes the system.

The errors on the values are quite small, but this is likely an artifact of the small parameter spaces and number of steps used during the MCMC. A higher number of steps and looser constraints would have been useful for helping the model to converge and for ensuring that the parameter space was sufficiently explored.

These inclination found by this model agrees with the inclinations of Icarus-A and Phoebe-B, but the mass ratio for this model doesn't agree with any of the other models. The mass of the neutron star agrees with the previous study and Icarus-A, but the companion mass doesn't agree with the other models' predictions.

3.10 Phoebe Model C - Constrained Fit with TESS Data

This final model is a replication of the Yi et al. (2022) study of J1123, with the only difference being that the companion mass was not fixed in the absence of a radial velocity curve. The priors used for this model are given in Table 3.2.

As with the B-models and other C-model, this model ran for 10,000 steps. Like the Icarus-C model, TESS data were used during the fitting procedure in lieu of the ZTF data.

The results from the fitting optimisation are given in Table 3.3. The inclination, mass ratio and both star masses all agree with those from the Yi et al. (2022) study, thus replicating their results.

3.10.1 Comparison with ZTF Data

As before, the ZTF data were also modelled with the best-fit parameters to investigate whether this fit can adequately describe the system. The light curves, and

residuals, are seen in Figure 3.9.




Figure 3.9: The top panels show the light curves across the ZTF g-band, r-band and i-band, and the TESS band with the *Phoebe-C* models fitted to the data. The bottom panels show the residuals for these bands.

There are obvious structures in the residuals, as was present in the previous study by Yi et al. (2022). Their light curve residuals also featured a periodic curve, indicating this model is a good replication of their model. The parameters are insufficient for both ZTF and TESS data, so it is understandable why the previous study resorted to more complex heating models though it is likely that looser constraints over a longer MCMC would have eliminated this structure - as was seen in *Icarus-A*.

The χ^2_{red} values for the different bands were the largest for each of the ZTF bands in comparison to the other five models, which support this fit not being adequate for describing the system. The TESS value was good, 1.20, but there is still a structure in those residuals. This model had the highest total χ^2_{red} value, at 13.23.

3.10.2 Summary

This model verified the results of the study by Yi et al. (2022) but indicates that the model is not sufficient at describing J1123. There is still prominent structure in the residuals.

The errors, though small, likely stem from the model have insufficient time (only 10,000 steps) to explore a small parameter space.

3.11 Comparison

The priors used by each model are summarised in Table 3.2, and the results are in Table 3.3. These were compared with the results of the previous study by Yi et al. (2022), which are summarised in Table 3.1.

The results of the previous study were successfully replicated by Phoebe-C, and other models agree with the previous study for some parameters but not all.

Phoebe-A became stuck in a local probability minimum during the MCMC process and so the results are significantly different to the other five models. Phoebe-A was restarted several times with varied constraints in an attempt to avoid the walkers becoming trapped in the local minimum, however the model consistently became stuck. Despite the results not being similar, the takeaway from this model is that Phoebe requires some extent of constraint in order to fit a model to the data.

The B and C models all indicate that the 10,000 step MCMC is insufficient for reaching a best-fit model that adequately describes the data. This, in turn, indicates that conclusions drawn in the previous study that suggested spot models were necessary to describe the system were not needed, rather the model needed to run for more steps across larger parameter spaces.

Icarus-A fared the best and though the errors on this model are larger than those from the previous study, the unconstrained run reached very similar conclusions to the heavily constrained Yi et al. (2022) run. Having run for sufficient time to converge, this model explored a large parameter space with minimal constraints. Some tighter constraints on the initial parameter spaces might reduce the errors from Icarus, but overall, this model outperforms the previous study model and indicates complex heating models are not necessary to describe J1123.

This indicates that Icarus is the preferable choice in the absence of additional fitting information, such as SED fitting, and in the presence of this additional information, either code will produce similar results though there may be some discrepancies.

Chapter 4

Discussion

Following the investigations of the "under-the-hood" behaviour of both codes, and comparison of their application to observational data, there is much to say about the codes. This chapter will discuss the results of this study and how what was found in Chapter 2 explains the results of Chapter 3.

4.1 Physical Accuracy

Phoebe accounts for a wide variety of physical effects, while Icarus makes more approximations to achieve similar - or even identical - results, in less time.

However, it is worth noting that the current signal-to-noise (SNR) ratios for spiders aren't very high (the current accuracy is not enough for the 2% discrepancy found in this study to be observationally detectable) do not cause Icarus to suffer against Phoebe, but with higher resolution standards expected from the Vera Rubin Observatory and the Roman Space Telescope in the coming years, Icarus may need to be upgraded in order to adequately account for all the physical affects present. Phoebe, on the other hand, can already account for time-dependent and relativistic effects.

In Chapter 2, it was clear that both Icarus and Phoebe use their photometry grids in similar ways and produce comparable results, though the difference in how the temperature, radius and surface gravity is calculated leads to small discrepancies in the simulated light curves of simulated systems. This doesn't cause

significant divergence, aside from for the most distorted part of the companion star.

From Chapter 2, the most prominent differences are clear. The proverbial "under the hood" behaviour of the codes, while mostly the same, differs in the calculations of the radius of the star at the Roche lobe distorted front of the star, which has a knock-on effect on surface gravity and temperature predictions across the surface, and in some differences in the temperature taken by both codes.

The Roche lobe filling factors can be made to match between both codes, producing near-identical light curves for the modelled systems, however the actual parameter value to generate the models differ. The same is true for the values for the temperature values being set.

The discrepancies that arise are focused at the front of the companion, leading to deviations in the produced light curves. This extent of these differences appears to decrease as the filling factor does - in J1123, the lower filling factor predicted by Phoebe in Yi et al. (2022) matches with the value predicted by Icarus-A.

In fact, all the discrepancies are focused at the low viewing angle, low temperature part of the companion: the front, facing the pulsar. This is seen in the variations in intensity with temperature and viewing angle in Chapter 2.

Accounting for these differences allows for light curves to match significantly better than when the filling factor and temperature input parameters are set to be identical. As such, it is expected that when real data is analysed, there will be a small discrepancy between the filling factor values and temperature values, but that as the Roche lobe filling factor decreases, so too will the discrepancy.

Only two of the models in Chapter 3 vary the filling factor and temperature. The previous study had a filling factor value of 0.90 ± 0.02 , which, when volume averaged, becomes 0.729 ± 0.050 , which the Icarus-A model matches with its value of $0.707^{+0.021}_{-0.009}$.

The other Icarus-A results were similar to those found in the Yi et al. (2022) study. The errors associated with this model were larger than those from the study, but the model was also unconstrained in the beginning whereas the previous study had significant constraints on their model. Smaller errors are seen in the B and C models for this reason.

The Phoebe-A results indicate a need for this code to have tighter constraints in the beginning of the fitting procedure to help the MCMC converge to more ap-

propriate parameter values. The higher errors in Icarus-A *agree* with this finding, and both suggest that there is a higher uncertainty associated with modelling systems without some preliminary constraints.

The Icarus and Phoebe B and C models have smaller errors, but required a number of assumptions to constrain them prior to fit optimization. These assumptions are derived from the same assumptions in the previous study of J1123 by Yi et al. (2022). However, some of these constraints are not appropriate for spider systems and may be adding tight constraints to parameters where they are not necessary. This suggests that while preliminary constraints are important, models can be over-constrained.

The optimal fit for the J1123 system was not the primary focus of this study, the focus instead was whether the modelling codes agreed when given the same data. Nevertheless, strong fits for J1123 were found and there is scope for optimizing the initial conditions necessary for both codes to ensure adequate exploration of parameter spaces without under or over constraining the models.

In terms of physical accuracy, there is no way to establish whether one code is more accurate than another, though the agreement of both the previous study and several of the Icarus and Phoebe models are strong indications of the true system parameters.

This suggests the codes can and should be used comparatively to strengthen confidence in parameter estimates for a system, though it is necessary to remember that they will have slightly different predictions as they do not create the same distortion on the companion star when given the same initial information.

There is one caveat to Phoebe’s physical accuracy that warrants mention: the method by which light curves are normalised may introduce biases. As Phoebe uses the median in its normalisation, if any portions of the data are missing or particularly noisy, the median value may be inaccurate, especially when the data are not uniformly sampled, leading the light curves and the subsequent models generated to be inaccurate. Issues like this did not arise in this study as a result of the quality of the data used, but data quality can vary.

4.2 Versatility

Phoebe is a code with applications beyond just spider binaries, so it accounts for a wide variety of different physical effects, which negatively impact its run-time. Icarus, on the other hand, is specifically for spiders. It has been designed explicitly to account for the dynamics of such systems.

Both codes are adaptable to various photometric grids, and both can implement a variety of complex heating models. Older editions of Phoebe can also account for relativistic effects, like Doppler boosting, though the version used in this study, Phoebe 2.4, no longer has this capability. Icarus also has Doppler boosting capabilities, when needed.

There is little between the most recent models of both codes to distinguish them in terms of the physics they can account for, in terms of spiders.

A notable difference arises in the constraints necessary to fit the model. Icarus is more versatile in that it reaches the same approximate conclusions with loose constraints as models from both codes with much stricter constraints. Phoebe, on the other hand, struggled to approach similar values to the other models when loosely constrained.

For systems where less information is known, or less can be inferred as was done in Yi et al. (2022) with SED fitting, then Icarus may be the more appropriate choice.

4.3 Speed

The most notable difference between the codes is how long it takes to generate a model. Icarus can generate a model in seconds, but Phoebe can take several minutes.

When fitting a model using the MCMC algorithm, this discrepancy is very apparent. To run a 100 step chain, with 50 walkers, it takes Icarus less than one minute but it would take Phoebe nearly an hour.

In situations where high speed computing is unavailable, as was the case for this study, using Phoebe becomes significantly more time consuming and, in the case of studies with a deadline, again as was the case for this study, Phoebe becomes a

risky choice. There may be ways to accelerate the rate at which Phoebe generates models but the documentation didn't indicate this.

While there is always a chance a code may fail to converge, Icarus' superior model computation speed is preferable. And as the parameter predictions for both codes are approximately the same based on the conclusions of this study, the approximations and interpolations used by Icarus in order to achieve this swiftness do not hinder its accuracy.

4.4 Overall

The parameters of a spider are not something that can be determined without the use of modelling codes, such as Icarus or Phoebe. Both, when given the same data, can reach similar conclusions about the parameter values that best describe the system, though they require different levels of constraints to do so.

The speed at which these codes can fit a model varies, with Icarus significantly faster than Phoebe without compromising its accuracy.

Discrepancies amounting up to $\approx 2\%$ arise between the models the codes generated from photometry grids, but current data resolutions are not sensitive enough for this to be an issue.

Modelling a spider system can be done successfully using either code, though they will likely not give identical results.

Chapter 5

Conclusions

The science upon which both codes, Icarus and Phoebe, are built is mostly identical. There are some crucial differences which affect the predictions the codes make somewhat. Differences in how the surface gravity is calculated, and deviations in the radius at the front of the Roche lobe distorted companion, cause discrepancies to arise in the surface gravities and temperatures, with the most noticeable differences occurring at the misshapen nose of the companion star, leading to small variations in predicted light curves.

The intensities are handled in similar ways, though stored in terms of different variables which requires some conversion when examining just the intensities, and the intensities vary across viewing angles and temperatures in both codes and the differences between them are small ($\approx 2\%$) enough to be undetectable with current data resolution limits.

The codes make similar parameter predictions when given the same data, though the amount of pre-required information differs. Icarus is more useful for systems with less constraints, while either code could be used for normalised data, though Phoebe may be a more straightforward option (as no edits will be necessary to account for the normalisation). There is the potential for some bias in Phoebe's normalisation process, resulting from the dependence on orbital sampling in the code, in which case, Icarus may be a more suitable choice.

The modelling of the spider candidate LAMOST J112306.9+400736, which had been previously modelled using Phoebe in Yi et al. (2022), verified informa-

tion known about the system and was an ideal test of how the codes worked with real data and what information was required for the modelling codes to converge to (or tend toward) optimised, best fit models. The mass values for the neutron star in this system verify the previous prediction, ranging between $1.003M_{\odot}$ and $1.414M_{\odot}$ (excluding the value predicted by Phoebe-A), agreeing with the value from the past study, $1.24M_{\odot}$.

The ultimate choice between both codes is personal to each researcher and their science goals. Either can be used effectively in modelling spiders, though in a direct head-to-head, Icarus would be a more versatile choice due to its superior speed and not needing as many constraints on its priors.

The codes are interchangeable and their results are comparable, though it will need to be noted that Phoebe requires much tighter priors than Icarus and this *may* affect some comparisons.

This study, spurred by research detailed in Sen et al. (2024), concludes that the binary modelling codes, Phoebe and Icarus, are quite similar in both how they work and what values they predict for observational data. There are deviations but as long as these are noted during comparisons between the codes, their results should be in agreement.

Bibliography

- Abbott, R., Abbott, T. D., Acernese, F., et al. 2023, *The Astrophysics Journal*, 955, 155
- Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, *Nature*, 300, 728
- Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, *Science*, 324, 1411
- Arzoumanian, Z., Nice, D. J., Taylor, J. H., & Thorsett, S. E. 1994, *The Astrophysics Journal*, 422, 671
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *The Astrophysics Journal*, 935, 167
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *The Astronomical Journal*, 156, 123
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *Astronomy & Astrophysics*, 558, A33
- Baade, W. & Zwicky, F. 1934, *Physical Review*, 46, 76
- Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, *Nature*, 300, 615
- Barr, E. D., Dutta, A., Freire, P. C. C., et al. 2024, *Science*, 383, 275–279
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 018002
- Bessell, M. & Murphy, S. 2012, *Publications of the Astronomical Society of the Pacific*, 124, 140–157

- Breton, R. P., Rappaport, S. A., van Kerkwijk, M. H., & Carter, J. A. 2012, *The Astrophysical Journal*, 748, 115
- Castelli, F. & Kurucz, R. L. 2003, in *IAU Symposium*, Vol. 210, *Modelling of Stellar Atmospheres*, ed. N. Piskunov, W. W. Weiss, & D. F. Gray, A20
- Chadwick, J. 1932, *Nature*, 129, 312
- Clark, C. J., Nieder, L., Voisin, G., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 502, 915–934
- Condon, J. J. & Ransom, S. M. 2016, *Essential Radio Astronomy*, sch - school edition edn. (Princeton University Press)
- Copperwheat, C. M., Marsh, T. R., Dhillon, V. S., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 402, 1824
- Cunningham, J.-M. C., Feliz, D. L., Dixon, D. M., et al. 2020, *The Astronomical Journal*, 160, 187
- Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, *The Astrophysical Journal*, 741, 103
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306
- Fruchter, A. S., Stinebring, D. R., & Taylor, J. H. 1988, *Nature*, 333, 237
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021, *A&A*, 650, C3
- Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, *The Astrophysical Journal*, 887, 93
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709
- Hogg, D., Bovy, J., & Lang, D. 2010, *ArXiv e-prints*
- Hui, C. Y. & Li, K. L. 2019, *Galaxies*, 7
- Jiang, L., Zhang, C.-M., Tanni, A., & Zhao, H.-H. 2013, in *International Journal of Modern Physics Conference Series*, Vol. 23, *International Journal of Modern Physics Conference Series*, 95–98

- Kennedy, M. R., Breton, R. P., Clark, C. J., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 3001
- Klinglesmith, D. A. & Sobieski, S. 1970, *The Astronomical Journal*, 75, 175
- Kreidberg, L., Bailyn, C. D., Farr, W. M., & Kalogera, V. 2012, *Astrophysics Journal*, 757, 36
- Kyer, R., Roy, S., Strader, J., et al. 2025
- LIGO Collaboration, T., Virgo Collaboration, t., KAGRA Collaboration, t., et al. 2024, Search for gravitational waves emitted from SN 2023ixf
- Linares, M. 2020, in *Multifrequency Behaviour of High Energy Cosmic Sources - XIII. 3-8 June 2019. Palermo*, 23
- Lorimer, D. R. & Kramer, M. 2004, *Handbook of Pulsar Astronomy*, Vol. 4
- Lucy, L. B. 1967, , 65, 89
- Ludlam, R., Cackett, E., Garcia, J., et al. 2022, in *AAS/High Energy Astrophysics Division*, Vol. 19, AAS/High Energy Astrophysics Division, 110.45
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *The Astronomical Journal*, 129, 1993–2006
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 018003
- Oppenheimer, J. R. & Volkoff, G. M. 1939, *Phys. Rev.*, 55, 374
- Orosz, J. A. & Hauschildt, P. H. 2000, The use of the NEXTGEN model atmospheres for cool giants in a light curve synthesis code
- Özel, F., Psaltis, D., Narayan, R., & McClintock, J. E. 2010, , 725, 1918
- Paczynski, B. 1971, , 9, 183
- Paerels, F., Mendez, M., Agueros, M., et al. 2009
- Prša, A., Conroy, K. E., Horvat, M., et al. 2016, *The Astrophysics Journals*, 227, 29

- Prša, A. & Zwitter, T. 2005, *The Astrophysics Journal*, 628, 426
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003
- Roberts, M. S. E. 2013, in *IAU Symposium*, Vol. 291, *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*, ed. J. van Leeuwen, 127–132
- Romani, R. W., Kandel, D., Filippenko, A. V., Brink, T. G., & Zheng, W. 2022, *The Astrophysical Journal Letters*, 934, L17
- Ruderman, M., Shaham, J., & Tavani, M. 1989, *Astrophysics Journal*, 336, 507
- Sen, B., Linares, M., Kennedy, M. R., et al. 2024, *The Astrophysical Journal*, 973, 121
- Strader, J., Swihart, S., Chomiuk, L., et al. 2019, *The Astrophysics Journal*, 872, 42
- Stringer, J. G., Breton, R. P., Clark, C. J., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 507, 2174–2191
- Tobin, R. W. & Berrington, R. C. 2024, *New Astronomy*, 109, 102210
- Vinciguerra, S., Salmi, T., Watts, A., et al. 2023, *An updated mass-radius analysis of the 2017-2018 NICER data set of PSR J0030+0451*
- Voisin, G., Kennedy, M. R., Breton, R. P., Clark, C. J., & Mata-Sánchez, D. 2020, *Monthly Notices of the Royal Astronomical Society*, 499, 1758
- von Zeipel, H. 1924, *Monthly Notices of the Royal Astronomical Society*, 84, 665
- Woosley, S. E., Heger, A., & Weaver, T. A. 2002, *Reviews of Modern Physics*, 74, 1015
- Yap, Y. X. J., Kong, A. K. H., & Li, K.-L. 2023, *The Astrophysics Journal*, 955, 21
- Yi, T., Gu, W.-M., Zhang, Z.-X., et al. 2022, *Nature Astronomy*, 6, 1203–1212

Appendix A

Parameter Relation Derivations

A.1 Mass Estimation Equations

The mass relations necessary for this project were not directly derived.

A.1.1 Icarus

The equation used to calculate the mass of the neutron star in Icarus is taken from Kennedy et al. (2022).

$$M_{NS} = \frac{Pk_{comp}^3(1+q)^2}{2\pi G \sin^3 i} \quad (\text{A.1})$$

where M_{NS} is the mass of the neutron star, P is the orbital period, k_{comp} is the radial velocity amplitude of the companion star, $q = \frac{M_{NS}}{M_{comp}}$ is the mass ratio, G is Newton's gravitational constant and i is the inclination.

A.1.2 Phoebe

The mass equation used in the Phoebe calculations is a minor adaption of that used in Icarus.

$$M_{NS} = \frac{4\pi^2 a^3}{GP^2(1+q)} \quad (\text{A.2})$$

where, again, M_{NS} is the mass of the neutron star, a is the semi-major axis of the binary, G is Newton's gravitational constant, P is the orbital period and $q = \frac{M_{NS}}{M_{comp}}$ is the mass ratio.

A.2 Semi-major Axis and Companion's Radial Velocity

To relate the available parameters in Phoebe with those in Icarus, the mass equations, A.1 and A.2 were equated.

$$\frac{4\pi^2 a^3}{GP^2(1+q)} = \frac{Pk_{comp}^3(1+q)^2}{2\pi G \sin^3 i} \quad (\text{A.3})$$

$$a^3 = \frac{Pk_{comp}^3(1+q)^2}{2\pi G \sin^3 i} \frac{GP^2(1+q)}{4\pi^2}$$

The gravitational constants cancel:

$$a^3 = \frac{Pk_{comp}^3(1+q)^2}{2\pi \sin^3 i} \frac{P^2(1+q)}{4\pi^2}$$

$$a^3 = \frac{P^3 k_{comp}^3 (1+q)^3}{8\pi^3 \sin^3 i}$$

$$a^3 = \frac{P^3 k_{comp}^3 (1+q)^3}{(2\pi)^3 \sin^3 i}$$

$$\sqrt[3]{a^3} = \sqrt[3]{\frac{P^3 k_{comp}^3 (1+q)^3}{(2\pi)^3 \sin^3 i}}$$

$$a = \frac{Pk_{comp}(1+q)}{(2\pi) \sin i} \quad (\text{A.4})$$

Appendix B

Corner Plots

For ease of reading, the corner plots for the six models from Chapter 3 have been included here.

A corner plot shows the relationship between parameters, and can be useful for

virtualising the distribution of the walkers in an MCMC parameter spaces.




Figure B.1: The corner plot for *Icarus-A*, showing the varied parameters - base temperature, inclination, filling factor, mass ratio, radial velocity amplitude, distance modulus - and inferred parameters - pulsar and companion star mass.

any discussion of this image?




Figure B.2: The corner plot for *Icarus-B*, showing the varied parameters - inclination and mass ratio - and inferred parameters - pulsar and companion star mass.




Figure B.3: The corner plot for *Icarus-C*, showing the varied parameters - inclination and mass ratio - and inferred parameters - pulsar and companion star mass.




Figure B.4: The corner plot for *Phoebe-A*, showing the varied parameters - effective temperature, inclination, equivalent radius, mass ratio, radial velocity amplitude.




Figure B.5: The corner plot for *Phoebe-B*, showing the varied parameters - inclination and mass ratio - and inferred parameters - pulsar and companion star mass.




Figure B.6: *The corner plot for Phoebe-C, showing the varied parameters - inclination and mass ratio - and inferred parameters - pulsar and companion star mass.*